
DYNAMICS OF QUANTUM VORTICES AT FINITE TEMPERATURE

Tod Martin Wright

A thesis submitted for the degree of
Doctor of Philosophy
at the University of Otago,
Dunedin, New Zealand.



March 2010

Copyright © 2010 Tod Martin Wright
All rights reserved.

This thesis is based on work published in references [267], [269], [270] and [268], copyright © 2008, 2009, 2010, and 2010 the American Physical Society, respectively.

Figure 4.1 has been adapted with permission from reference [28], copyright © 2008 Taylor & Francis.

Version history

Original version submitted for examination: March 26, 2010.

Revised version post-examination: October 26, 2010.

Version submitted to arXiv: November 26, 2010.

Abstract

We perform investigations into the behaviour of finite-temperature degenerate Bose gases using a classical-field formalism. Applying the projected Gross-Pitaevskii equation approach to the description of equilibrium and nonequilibrium regimes of the degenerate Bose gas, we consider the interpretation of the classical-field trajectories, and their relation to mean-field descriptions of finite-temperature Bose condensates. We show that the coherence of the field can be characterised by its temporal correlations, and discuss how the symmetry-broken averages familiar from mean-field theories emerge from the classical-field trajectories.

The major focus of this thesis is the dynamics of quantum vortices in finite-temperature Bose fields. We show that a finite-temperature condensate containing a precessing vortex in a cylindrically symmetric trap can be realised as an ergodic equilibrium of the classical-field theory. We demonstrate the identification of the rotationally symmetry-broken condensate orbital from the short-time correlations of the field, and show how the thermal core-filling of the vortex emerges from the classical-field fluctuations. We then consider the nonequilibrium dynamics that result when such a precessing-vortex configuration is subjected to a static trap anisotropy which arrests its rotation. We decompose the nonequilibrium field into condensed and noncondensed components, and observe novel coupled relaxation dynamics of the two components.

Finally, we consider the nucleation of vortices in an initially zero-temperature quasi-two-dimensional condensate stirred by a rotating trap anisotropy. We show that vacuum fluctuations in the initial state provide an irreducible mechanism for breaking the initial symmetries of the condensate and seeding the subsequent dynamical instability. A rotating thermal component of the field quickly develops and drives the growth of oscillations of the condensate surface leading to the nucleation of vortices. We study the relaxation and rotational equilibration of the initially turbulent collection of vortices, and monitor the effective thermodynamic parameters of the field during the relaxation process. We show that the equilibrium temperature of the field is well predicted by simple arguments based on the conservation of energy, and that thermal fluctuations of the field prevent the vortices from settling into a rigid crystalline lattice in this reduced dimensionality. We find that true condensation in the field is completely destroyed by the disordered motion of vortices, but show that the temporal correlations of the field distinguish the quasi-coherent *vortex liquid* phase in the trap centre from the truly thermal material in its periphery.

Acknowledgements

I would like to express my gratitude to my supervisor, Professor Rob Ballagh. Rob is an enthusiastic physicist, a dedicated mentor, the consummate professional, and above all a genuinely nice guy. I have benefited greatly from his wisdom, guidance, and support over the last few years. Despite his ever-growing workload, Rob has always found time to discuss my research, offer professional advice, and even just chat, for which I am very grateful. I thank him for everything he has taught me, and everything he allowed me to learn for myself.

My co-supervisor Assoc. Prof. P. Blair Blakie has always made himself available to give advice and help where he can, for which I am also very grateful.

I would particularly like to thank Dr. Ashton Bradley, not only for giving me the code employed in much of the research of this thesis, but also for his input into and support of my research, initially from a distance, but with increased patience and dedication after his move to Otago. I would also like to thank Prof. Crispin Gardiner for his input into my work, his encouragement, and his patronage.

I have learned a lot in my brief acquaintance with Assoc. Prof. Colin Fox. Colin has a staggering knowledge and understanding of a wide range of matters physical, mathematical, and statistical. I thank him for his assistance with mathematics, his miscellaneous wisdoms and insights, and for broadening my perspective. I would also like to thank Terry Cole and (more recently) Peter Simpson for their help with computing, and for everything I have learned from them.

Special thanks are due to Assoc. Prof. Lars Bojer Madsen, Prof. Klaus Mølmer, Dr. Nicolai Nygaard, Dr. Katherine Challis, and others at Aarhus University, for their hospitality during my brief visit to Denmark. I'm also grateful to (among others) Dr. Andrew Sykes, Assoc. Prof. Matthew Davis, Dr. Thomas Hanna, Dr. Simon Gardiner, and Dr. Robin Scott for being familiar faces at conferences here and abroad, and for taking the time to talk physics with me.

I would like to thank all the other students that have shared the postgraduate experience with me, and in particular Bryan Wild, Emése Tóth, and Dr. Alice Bezett, who have been supportive officemates throughout the majority of my time in room 524b. Sam Lowrey deserves special mention for frequently pulling me away from my work to get some much-needed fresh air and perspective. I am grateful also to my friends outside physics,

particularly to Luke, Morgan, and Anna, for regularly taking me away from it all, and for understanding when this hasn't been feasible.

I am grateful to my sisters Milvia and Layla, for their support and encouragement. Finally, I would like to thank my parents, for their belief in me, their support, and their encouragement, and in particular my father, for sparking my interest in science at an early age.

This work was supported by a University of Otago Postgraduate Scholarship; the Marsden Fund of New Zealand under Contract No. UOO509; and the New Zealand Foundation for Research, Science and Technology under Contract No. NERF-UOOX0703.

Contents

1	Introduction	1
1.1	Ultracold atoms and Bose-Einstein condensation	1
1.2	Quantum vortices	3
1.3	This work	5
1.3.1	Thesis overview	5
1.3.2	Peer-reviewed publications	7
2	Background theory	9
2.1	Identical-particle ensembles	9
2.1.1	Second quantisation	10
2.2	Degenerate bosons	12
2.2.1	Bose condensation	12
2.2.2	The Gross-Pitaevskii equation	15
2.2.3	Bogoliubov theory	21
2.3	Superfluidity and vortices	28
2.3.1	Superfluidity	28
2.3.2	Vortices	30
2.3.3	Vortices in Bose condensates	33
3	Classical-field theory	39
3.1	Nature of the classical-field approximation	40
3.2	Projection	41
3.2.1	Motivation	41
3.2.2	Projected classical-field theory	45
3.3	Ergodicity	49
3.3.1	Correlation functions	50
3.3.2	Thermodynamics	52
3.4	Truncated-Wigner approximation	54

Contents

3.4.1	The Wigner distribution	55
3.4.2	Truncated-Wigner approximation for the interacting Bose gas . . .	58
3.4.3	Relevance to the PGPE formalism	60
4	Numerical methods	63
4.1	Numerical integration of the projected Gross-Pitaevskii equation	63
4.1.1	Dimensionless formalism	63
4.1.2	Spectral-Galerkin formulation	66
4.1.3	Interaction-picture methods	67
4.1.4	Gaussian-quadrature evaluation of projected operators	68
4.2	Performance of the Gauss-Laguerre algorithm	70
4.2.1	Base rotating-frame algorithm	70
4.2.2	Perturbing potential term	72
4.3	Comparison to other discretisation schemes	73
4.3.1	Phase-space covering	74
4.3.2	Implementing spatial derivatives	77
5	Anomalous moments in classical-field theory	83
5.1	System	84
5.2	Temporal coherence and the emergence of a nonzero first moment	85
5.2.1	Identification of the first moment	85
5.2.2	Temporal coherence and sample length	88
5.3	Dependence of the first moment on the field energy	93
5.3.1	Condensate fraction	93
5.3.2	Condensate frequency	95
5.4	Pairing correlations	96
5.4.1	Methodology	98
5.4.2	Anomalous density	99
5.4.3	Dependence on field energy	101
5.5	Summary	102
6	Equilibrium precession of a single vortex	105
6.1	Simulation procedure	106
6.1.1	System	106
6.1.2	Microcanonical evolution	107
6.2	Equilibrium behaviour of the field	110
6.2.1	Vortex precession	112
6.3	Field-covariance analysis	114

6.3.1	Lab-frame analysis	115
6.3.2	Rotating frames	119
6.3.3	Temporal decoherence and sample length	123
6.4	Rotational properties of the field	126
6.4.1	Condensate and noncondensate observables	128
6.4.2	Angular momentum of the condensate	129
6.4.3	Angular velocity of the thermal cloud	129
6.5	Dependence on internal energy	130
6.5.1	Density	130
6.5.2	Condensate fraction and angular momentum	130
6.5.3	Condensate and cloud rotation rates	131
6.5.4	Temperature and chemical potential	134
6.5.5	Moment of inertia and vortex displacement	137
6.6	Dependence on frame of projector	138
6.7	Summary	140
7	Nonequilibrium dynamics of vortex arrest	143
7.1	Simulation procedure	144
7.1.1	Initial state	144
7.1.2	Evolution	145
7.2	Arrest dynamics	145
7.2.1	Vortex trajectory	145
7.2.2	Rotational dynamics of condensate and cloud	147
7.2.3	Heating of the atomic field	151
7.3	Dependence on trap anisotropy	154
7.3.1	Weak damping	154
7.3.2	Strong damping	156
7.3.3	Decay times	157
7.4	Summary	158
8	Prior studies of vortex-lattice formation in stirred condensates	161
8.1	Experiments	162
8.1.1	Work of Dalibard's group (ENS)	162
8.1.2	Work of Ketterle's group (MIT)	163
8.1.3	Work of Foot's group (Oxford)	163
8.2	Theory	164
8.3	Outstanding questions	171

9	Condensate stirring	175
9.1	System and simulation procedure	176
9.1.1	The system	176
9.1.2	Initial state: inclusion of vacuum noise	176
9.1.3	Evolution	179
9.2	Simulation results	180
9.3	Analysis	183
9.3.1	Thermodynamic parameters	184
9.3.2	Rotational parameters	187
9.3.3	Temporal analysis	191
9.3.4	Vortex motion	202
9.4	Dependence on simulation parameters	204
9.4.1	Chemical potential	205
9.4.2	Cutoff dependence	211
9.5	Summary	214
10	Conclusion	217
10.1	Thesis summary	217
10.2	Directions for future work	219
A	Condensation in the noninteracting classical field	221
B	Interaction-picture integration algorithms	223
B.1	The RK4IP algorithm	223
B.1.1	Fourth-order Runge Kutta	224
B.1.2	RK4IP	224
B.2	The ARK45-IP algorithm	225
B.2.1	Adaptive Runge-Kutta	226
B.2.2	ARK45-IP	227
C	Quadrature methods for the PGPE	229
C.1	Gaussian quadrature rules	229
C.2	Hermite quadratures for harmonic confinement	230
C.2.1	Application to the PGPE	230
C.3	Laguerre quadratures for the rotating system	233
C.3.1	Nonlinear term	233
C.3.2	Elliptic trap anisotropy	236

D Semiclassical fitting function	239
References	243

Chapter 1

Introduction

1.1 Ultracold atoms and Bose-Einstein condensation

Since the first experimental achievement of *Bose-Einstein condensation* in dilute alkali gases in 1995 [8, 62, 38], research into Bose condensates and other quantum degenerate phases of matter in cold gaseous samples of atoms has grown explosively. The field of *ultracold atomic physics* has become a unique nexus between researchers from a diverse array of subfields of modern physics, ranging from atomic physics, laser physics and quantum optics, to nuclear and condensed-matter physics, quantum information, fluid dynamics and even astrophysics and cosmology. In recent years the confinement and manipulation of ultracold atoms has expanded to encompass the study of many different atomic species (both bosonic and *fermionic* [69]), in a range of geometries [78, 185, 123], with a variety of different interatomic interactions [53, 154]. All of these studies share the common feature of investigating interacting many-body systems and their behaviour in a clean and precisely controlled environment [32].

Bose-Einstein condensates were historically the first ultracold atomic systems realised experimentally, and these systems are still intensively studied in laboratories throughout the world. A Bose condensate forms when a macroscopic number of identical *bosonic* atoms congregate into a single spatial mode. The remarkable aspect of the phase transition to the Bose-condensed state is that it is driven by the quantum statistics of the particles, rather than the interactions between them, and may thus occur even in a noninteracting sample of bosonic atoms. The phenomenon of Bose-Einstein condensation of massive particles was predicted in 1925 by Einstein, and London conjectured in 1938 that Bose condensation was responsible for the remarkable properties of the then recently discovered superfluid helium. However, the advent of *laser cooling* allowed researchers for the first time to observe the Bose condensation phenomenon in isolation, with the unmatched

precision of control and observation developed over many decades in atomic and laser physics.

As they are essentially composed of many atoms residing in a single spatial mode, Bose condensates provide an amplification of the underlying quantum mechanical physics of microscopic systems to an observable scale [164]. Condensates are characterised by their long-range quantum coherence, and form a close analogue of coherent laser light, despite being composed of massive, interacting particles. The quantum coherence properties of atomic Bose condensates have been demonstrated dramatically in experiments of matter-wave interference [10], mixing [70], and amplification [137].

An important property of Bose condensates is their *superfluidity*: like superfluid helium, they can support persistent currents and vortices. However, in contrast to the superfluid *liquid* helium, in which only $\sim 10\%$ of the atoms form a Bose condensate due to the strong depleting effect of interatomic interactions [201], the weakly interacting nature of the dilute gases allows the formation of essentially pure condensates. Moreover, the finite-sized and *inhomogeneous* nature of these gaseous superfluids leads to new features, and behaviour quite different to the more traditional, essentially uniform superfluids [114]. There has thus been intensive study of the thermodynamics of the trapped Bose gases, and of the critical regime associated with the Bose-condensation phase transition [77, 108]. Furthermore, additional phases of degenerate Bose gases have been studied in novel trapping geometries such as *optical lattices* [185], and in reduced dimensionalities, where new quasi-coherent phases of the gas appear [247, 78, 55].

The precision of control in atomic physics experiments also allows for the thorough investigation of the *dynamic* properties of dilute condensates. Experiments have considered not only the collective response of condensates to perturbing potentials [143, 178], but also the formation and dynamics of topological defects such as solitons [41, 147] and vortices [180, 175] in these quantum fluids. Experiments have even extended to consider the strongly *nonequilibrium* dynamics of these systems, including turbulent regimes of the atomic field [127] and nonequilibrium phase transitions [260].

In addition to the precision with which condensates can be observed and manipulated experimentally, their *weakly* interacting nature makes them an attractive system for theoretical investigation. In comparison to the enormous complexity of dense liquids such as superfluid helium, these dilute gases are amenable to *tractable* theoretical approaches, including essentially *ab initio* numerical calculations and, in many cases, analytical calculations. In particular, some of the theoretical tools [33] which could only give a qualitative level of insight into the strongly interacting superfluid helium found new life as practically *exact* descriptions of the dilute-gas condensates in many scenarios. The dilute condensates thus offer us an opportunity to understand superfluidity in terms of *microscopic* models of

the relevant physics. In addition to the traditional tools of superfluidity, condensed matter and nuclear physics, the close analogy between Bose condensates and coherent light fields has lead to the broad application of methods from *quantum optics* and laser physics to the description of condensates [28]. In general, however, many questions remain as to how to properly describe the dynamics of Bose-Einstein condensates in general situations, in which the atomic cloud possesses a significant thermal component, and is possibly far from equilibrium [217]. Moreover, heavy numerical calculations are usually required to understand the properties and dynamics of Bose condensates in detail. These numerical studies typically involve the solution of *nonlinear* equations, and so small inaccuracies and errors in their implementation can easily lead to completely unphysical results. It is thus of great importance not only to develop accurate theoretical models of Bose condensation, but also to construct accurate *numerical* models and methodologies for the simulation of condensates and their dynamics.

1.2 Quantum vortices

The ‘smoking gun’ of the superfluid transition in quantum fluids is the appearance of quantum *vortices*: tiny whirlpools in which the superfluid circulates about a ‘core’ where the superfluid density vanishes [75]. These vortex structures are topological defects, which reflect the inherent quantum-mechanical nature of a superfluid: due to the single-valued nature of the wave-function phase which determines the superfluid flow, the phase around a vortex must change by an exact multiple of 2π , and thus the rotation of the superfluid is fundamentally *quantised*. Singly charged vortices (with phase circulation $\pm 2\pi$) are topologically stable, in the sense that they must encounter another such defect (of opposite rotation) or a boundary of the fluid in order to decay. Vortices with higher charges are energetically unstable to decay by fissioning into two or more vortices of lesser topological ‘charge’ [75, 219]. Under rotation, superfluid helium mimics the velocity profile of a rotating rigid body by admitting an ensemble of many such singly charged vortices, which, due to their mutual repulsion [243], form a regular crystalline *lattice* at equilibrium.

The physics of superfluid vortices are of great general interest, as analogous objects occur in a variety of physical systems including type-II superconductors [31]. Above a critical magnetic imposed magnetic field strength, such superconductors admit thin filaments of magnetic flux, about which the charged ‘condensate’ of Cooper-paired electrons forms vortices. Vortex structures are also thought to occur in the dense nucleon-superfluid phases of rotating neutron stars [20]. In both of these systems the behaviour, and in particular the *dynamics*, of vortices are crucial to understanding the physics of the quantum

fluids involved. The nonequilibrium dynamics of *turbulence* in superfluids are of great interest due to the role of quantised vortices in such regimes. Indeed, the turbulence of *classical* fluids is significantly complicated by the continuous nature of vorticity in the system. Studies of so-called *quantum* or *superfluid* turbulence [19] are thus actively pursued in the hope that the study of turbulence in systems with quantised vorticity will yield insights into the classical turbulence problem.

The realisation of Bose-Einstein condensates and other degenerate superfluid phases of matter in atomic-physics experiments has provided new opportunities to study the intrinsic phenomenology of quantum vortices. The first vortices in Bose condensates were formed by coherent interconversion between two hyperfine components of a condensed atomic cloud, creating a structure in which one component circulates around a non-circulating core composed of the other component [180]. By selectively removing the core-filling component with resonant light pressure, an isolated vortex was obtained [7]. In further experiments, inspired by the *rotating-bucket* experiments of superfluid helium, atomic gases were evaporatively cooled to condensation in a rotating anisotropic trapping potential [175], producing both single vortices and small vortex arrays. This was followed by investigations of the dynamical stirring of a preformed, vortex-free condensate with a rotating potential anisotropy [176, 2, 130]. In this approach an initially nonrotating condensate could be transformed by a strongly nonequilibrium process into a rotating vortex lattice. Other experiments considered the formation of a vortex lattice by cooling a rotating thermal cloud through the Bose-condensation transition in a *cylindrically symmetric* trap, in order to observe the *intrinsic* dynamics of vortex nucleation driven by the rotation of the normal fluid [124]. The dynamics of vortex-lattice formation in this case are very rich, as the growth of the condensate and the formation of the vortex lattice are essentially intertwined nonequilibrium phase transitions [37]. Vortices have also been imprinted in condensates using topological (Berry) phases [159], and by ‘sweeping’ a laser light potential through the condensed atomic cloud [136]. Another scheme for forming a vortex involves the transfer of orbital angular momentum from a laser beam to the Bose condensate [6]. Vortices can form spontaneously in the merging of multiple Bose condensates, due to the essentially random relative phase difference that develops between the condensates [226], and can also be ‘trapped’ in the nascent condensate during its rapid growth from a (nonrotating) quenched thermal cloud [260]. These formation processes are closely related to conjectured mechanisms of phase-defect trapping in early-universe cosmology [148, 277].

Experiments have also investigated the role of vortices in the two-dimensional superfluid transition [247], and the effect of artificial pinning potentials on vortices [255]. The realisation of a *strongly interacting* superfluid in a degenerate *Fermi* gas was con-

clusively demonstrated by the observation of vortices in this system [278]. Rotating Bose condensates containing vortices have also been proposed as a platform for investigating fundamental many-body physics, including aspects of symmetry-breaking in many-body systems [58], and the potential realisation of *strongly correlated* states of vortex matter, such as bosonic analogues of the fractional-quantum-Hall states in rapidly rotating condensates [256]. The difficulty in attaining the extreme experimental conditions necessary to observe such states has spurred researchers to consider alternative scenarios involving (e.g.) *artificial gauge fields* [167] in order to obtain high vortex densities in the hope of observing the *quantum melting* of dense vortex lattices.

The weakly interacting nature of dilute atomic gases offers us hope of obtaining *quantitative* theoretical descriptions of vortex dynamics in atomic Bose condensates. However, the results of experiments involving strongly nonequilibrium dynamics of vortex formation and motion have often proved difficult to describe theoretically. Conspicuously, the standard mean-field approximations for Bose condensates are inapplicable in the turbulent regimes of Bose-field dynamics arising in vortex-formation scenarios, in which a clear distinction between condensed and noncondensed material cannot be made [170]. It is therefore important to develop dynamical theories of the degenerate Bose gas which can accurately describe the nonequilibrium dynamics of vortices, their interactions with the thermal component of the field, and turbulent behaviour in the degenerate Bose gas.

1.3 This work

1.3.1 Thesis overview

In chapter 2 we give the general theoretical background for this thesis. We begin by briefly reviewing the description of quantum many-body systems in first and second quantisation. We introduce the concept of Bose-Einstein condensation, first for noninteracting systems and then for interacting systems, and derive the Gross-Pitaevskii equation for the condensate wave function and the Bogoliubov theory for its excitations. We discuss the notion of superfluidity, and its relation to Bose-Einstein condensation. We then consider the quantum vortices which are a characteristic feature of superfluidity, and discuss their response to imposed forces, including the frictional effects of a normal (non-superfluid) component of the fluid.

In chapter 3 we introduce the classical-field methodology we employ throughout this thesis. We begin with a heuristic explanation of the fundamental classical-field approximation, before deriving the *projected Gross-Pitaevskii equation* for the evolution of classical-field trajectories. We describe how the *ergodicity* of this nonlinear equation of motion

for the classical atomic field gives rise to a practical means for evaluating correlations in the field, and discuss the emergence of thermodynamic behaviour from the fundamentally dynamical field equation. We conclude this chapter with a brief review of the concept of quasiprobability distributions and phase-space methods. Focussing in particular on the *truncated-Wigner* formalism, we relate the equation of motion obtained in the truncated-Wigner approach to the projected Gross-Pitaevskii equation, and discuss the implications of the truncated-Wigner formulation for our classical-field investigations in this thesis.

In chapter 4 we describe the techniques we use for the numerical integration of the projected Gross-Pitaevskii equation, and quantify the performance of the Gauss-Laguerre-quadrature algorithm we use for the simulation of rotating systems and vortices in this thesis. We then contrast the harmonic-basis projected Gross-Pitaevskii approach we adopt in this thesis to other common discretisations of the continuum Gross-Pitaevskii equation.

In chapter 5 we consider the temporal correlations that emerge in the classical-field description of a harmonically trapped Bose gas, and show how the condensate can be identified by its quasi-uniform phase rotation. We show that this provides an analogue of the symmetry-breaking assumption commonly employed in traditional (mean-field) theories of Bose condensation, and we identify the condensate as the anomalous first moment of the classical field. We then demonstrate the generality of this prescription by calculating the pair matrix which describes anomalous correlations in the thermal field, and calculate the anomalous thermal density of the field, which we find to have temperature-dependent behaviour consistent with that predicted by mean-field theories.

In chapter 6 we consider the equilibrium precession of a single vortex in a finite-temperature Bose-Einstein condensate. We show that a precessing-vortex configuration can be obtained as an ergodic equilibrium of the classical field with finite conserved angular momentum. The symmetry-breaking nature of this state requires us to go beyond formal ergodic averaging in order to characterise condensation in the field, and we show that an analysis of the short-time fluctuation statistics of the field reveals the symmetry-broken condensate orbital, the thermal filling of the vortex core, and the Goldstone mode associated with the broken rotational symmetry.

In chapter 7 we consider a nonequilibrium scenario of vortex arrest. We form an equilibrium finite-temperature precessing-vortex configuration in a cylindrically symmetric trap, and then introduce a static trap anisotropy. The anisotropy leads to the loss of angular momentum from the thermal cloud and, consequently, to the decay of the vortex. Our approach provides an intrinsic description of the coupled relaxation dynamics of the nonequilibrium condensate and thermal cloud. By considering the short-time correlations of the nonequilibrium field, we separate it into condensed and noncondensed components, and identify novel dynamics of angular-momentum exchange between the two compo-

nents during the arrest. We also discuss qualitatively distinct regimes of relaxation that we obtain by varying the strength of the arresting trap anisotropy.

The remainder of this thesis is concerned with the nucleation of quantum vortices by the stirring of a condensate with a rotating trap anisotropy. We begin in chapter 8 by briefly reviewing the prior literature on this topic. We discuss the experimental approaches to condensate stirring taken by research groups at École Normale Supérieure, MIT and Oxford. We then discuss the various theoretical and numerical approaches previously taken to describe the condensate-stirring process, and the more surprising features of the experimental results. We conclude this chapter by outlining the remaining open questions and controversy as to the mechanisms responsible for vortex nucleation and damping in these experiments.

In chapter 9 we present our classical-field approach to the stirring problem. We consider the stirring of a condensate initially at zero temperature in a highly oblate (quasi-two-dimensional) trapping potential. We add a representation of vacuum noise to the initial field configuration according to the truncated-Wigner prescription. This noise seeds the spontaneous processes responsible for the ejection of material from the unstable condensate, and we show that the ejected material quickly forms a rotating thermal cloud which drives the nucleation of vortices at the condensate surface. We quantify the development of the thermodynamic parameters of the field during its evolution, and show that the final temperature attained is well predicted by simple energy-conservation arguments. We find that the equilibrium state of the atomic field in this quasi-2D scenario is a disordered vortex liquid, in which vortices are prohibited from crystallising into a regular lattice by thermal fluctuations. We find that condensation in the field is destroyed by the vortex motion, but show that the temporal correlations in the field allow us to distinguish the quasi-coherent structure at the centre of the trap from the truly thermal material in its periphery.

Finally, a summary and conclusion is given in chapter 10, in which we outline possible directions for future research.

1.3.2 Peer-reviewed publications

Much of the work contained in this thesis has been published in refereed journals. The work on temporal coherence and anomalous averages in classical-field simulations presented in chapter 5 has appeared in *Physical Review A* [268]. The work on the equilibrium precession of a vortex and on the arrest of vortex rotation presented in chapters 6 and 7 has also been published in *Physical Review A* (references [269] and [270] respectively). The discussion and simulations of the condensate-stirring dynamics presented in chap-

Chapter 1. Introduction

ters 8 and 9 constitute an expanded and updated account of work also published in *Physical Review A* [267].

Chapter 2

Background theory

In this chapter we will review the general theoretical background for this thesis. The key results of many-body theory, including the formalism of second quantisation, are reviewed in many standard texts on the subject, for example references [88, 5, 25]. Here, we give a brief review of these topics, to illustrate the fundamental complexity of many-body systems such as the degenerate gases we seek to describe in this thesis, and to give context to the *classical-field* approximation to the many-body theory we introduce in chapter 3. We then proceed to review some standard definitions and theoretical tools for the study of Bose-Einstein condensation. Many excellent reviews of these topics exist in the literature (see references [164, 60, 146, 86, 207, 204, 47]), and hence we will review only a small selection of topics relevant for our discussions later in this thesis. Finally, we will discuss aspects of superfluidity and vortices relevant to our investigations in this thesis, primarily following material presented in references [86, 75, 87].

2.1 Identical-particle ensembles

In the study of *dilute* atomic gases, the atoms can usually be considered as an ensemble of N identical structureless particles which interact via a pair-wise (binary) potential $V(\mathbf{x}_1 - \mathbf{x}_2)$. In the fundamental or *first-quantised* description of such an assembly of particles in quantum mechanics, the dynamics of the ensemble are governed by the many-body Hamiltonian

$$H_{\text{MB}} = \sum_i H_{\text{sp}}(\mathbf{x}_i) + \sum_{i < j} V(\mathbf{x}_i - \mathbf{x}_j), \quad (2.1)$$

where the single-particle Hamiltonian

$$H_{\text{sp}}(\mathbf{x}) = \frac{-\hbar^2 \nabla^2}{2m} + V_{\text{ext}}(\mathbf{x}), \quad (2.2)$$

includes the kinetic energy of each atom and the effect of an *external* potential $V_{\text{ext}}(\mathbf{x})$. The state of the atoms is described by a many-body state in the appropriate many-particle Hilbert space, which is formed as the tensor product of the N individual single-particle Hilbert spaces [25]

$$\mathcal{H}_N = \mathcal{H}_1^{(1)} \otimes \mathcal{H}_1^{(2)} \otimes \cdots \otimes \mathcal{H}_1^{(N)}. \quad (2.3)$$

More specifically, the particular many-body state is required to exhibit a specific symmetry under permutations of the particle labels, depending on the spin of the particles involved. Formally, this requirement can be deduced for elementary particles from the considerations of relativistic quantum mechanics [198]. The atoms we consider here are of course composite particles, but they may be treated simply as structureless particles with spin given by the total spin of their composite nucleons and electrons, provided that (as is always the case in this thesis) the energy scales of their interaction do not probe their true internal structure [146]. For bosons (with integer spin in units of the reduced Planck's constant \hbar) the state must be symmetric with respect to the permutation of particle labels, and thus resides in the subspace \mathcal{H}_N^S of \mathcal{H}_N spanned by completely symmetric states, whereas in the case of Fermions (with half-integer spin) the state must be antisymmetric with respect to the permutation of labels (i.e., $\langle \cdots \mathbf{x}_j, \mathbf{x}_i, \cdots | \psi \rangle = -\langle \cdots \mathbf{x}_i, \mathbf{x}_j, \cdots | \psi \rangle$), and is thus a vector in the subspace \mathcal{H}_N^A spanned by completely antisymmetric states [25].

2.1.1 Second quantisation

We now discuss an alternative representation of the interacting particle ensemble [88] which is known as the *second-quantised* representation of the many-body system. This formulation promotes the discussion to that of a *quantum field theory*. The concept of a quantum field theory originates in relativistic quantum mechanics, in which it is found that the quantum theory is *necessarily* many-body in nature, as a consequence of mass-energy equivalence [261]. In the ultracold scenarios we consider in this thesis, the rest-mass of atoms dwarfs all other energy scales in the system, and atom-number changing processes do not occur. However, in studies of non-relativistic many-body systems, the second-quantised formalism has great advantages over the first-quantised formalism. First, it provides a convenient encapsulation of the wave-function symmetry constraints, which rapidly become unwieldy in describing ensembles of increasing particle number. Second, it is a convenient formulation in which to identify approximations to the (generally intractable) complete many-body description. Although the most widespread approximations schemes are based on the techniques of diagrammatic perturbation theory [88, 5], in this thesis we utilise a different approximation scheme: the second-quantised formulation of the many-body system allows us to identify a *classical limit*. Indeed in the case

of bosons, the appropriate quantum field theory can be obtained by building quantum fluctuations into the dynamics of a classical field which evolves in a *single-particle* space. Significant insight into the dynamics of the quantum many-body system can thus, in certain regimes, be obtained by studying this dramatically simpler classical field, as we discuss in chapter 3.

To formulate the second-quantised theory, we introduce the *field operator* $\hat{\Psi}(\mathbf{x})$, which obeys the relations

$$\left[\hat{\Psi}(\mathbf{x}), \hat{\Psi}(\mathbf{x}')\right]_{\pm} = \delta(\mathbf{x} - \mathbf{x}'); \quad \left[\hat{\Psi}(\mathbf{x}), \hat{\Psi}(\mathbf{x}')\right]_{\pm} = \left[\hat{\Psi}^{\dagger}(\mathbf{x}), \hat{\Psi}^{\dagger}(\mathbf{x}')\right]_{\pm} = 0, \quad (2.4)$$

where $[\cdot, \cdot]_{-}$ denotes a commutator, and applies in the case of bosons, and $[\cdot, \cdot]_{+}$ denotes an anti-commutator, which applies in the case of fermions. In this description, the operator $\hat{\Psi}(\mathbf{x})$ and its (Hermitian) conjugate are the fundamental operator-valued quantities, which act on the *Fock space* formed as the union of many-body spaces with all (physical) particle numbers [25]

$$\mathcal{F} = \mathcal{H}_0 \oplus \mathcal{H}_1 \oplus \mathcal{H}_2 \oplus \cdots. \quad (2.5)$$

It is important to note that the coordinates \mathbf{x} in this formalism are not operators, and simply label the field operators. In this formulation the (first-quantised) many-body Hamiltonian (2.1) can be rewritten *equivalently* as the second-quantised Hamiltonian

$$\hat{H} = \int d\mathbf{x} \hat{\Psi}^{\dagger}(\mathbf{x}) H_{\text{sp}} \hat{\Psi}(\mathbf{x}) + \frac{1}{2} \int d\mathbf{x} \int d\mathbf{x}' \hat{\Psi}^{\dagger}(\mathbf{x}) \hat{\Psi}^{\dagger}(\mathbf{x}') V(\mathbf{x} - \mathbf{x}') \hat{\Psi}(\mathbf{x}') \hat{\Psi}(\mathbf{x}). \quad (2.6)$$

A general single-particle operator on the many-body space in the first-quantised treatment acts on *all* the indistinguishable particles, and is thus of form

$$J = \sum_{i=1}^N J(\mathbf{x}_i), \quad (2.7)$$

where $J(\mathbf{x}_i)$ acts on the i^{th} particle. Such an operator is promoted to an operator on the Fock space by the correspondence [85]

$$J \rightarrow \hat{J} = \int d\mathbf{x} \hat{\Psi}^{\dagger}(\mathbf{x}) J(\mathbf{x}) \hat{\Psi}(\mathbf{x}), \quad (2.8)$$

and two- (and higher) body operators in the first-quantised formalism can be promoted to operators on the Fock space by similar correspondences. We stress that if the system is in a state $|\chi\rangle$ (on the Fock space) with *definite* particle number N , then this formulation is *equivalent* to that of the first-quantised many-body theory, with the many-body wave

function given by [229]

$$\psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \frac{1}{\sqrt{N!}} \langle \text{vac} | \hat{\Psi}(\mathbf{x}_N) \cdots \hat{\Psi}(\mathbf{x}_2) \hat{\Psi}(\mathbf{x}_1) | \chi \rangle, \quad (2.9)$$

where $\langle \text{vac} |$ is the bra corresponding to the state with no particles (the vacuum). However, the second-quantised formalism also allows us to describe states with *indefinite* particle numbers. In the remainder of this thesis, we will restrict our attention to bosons.

2.2 Degenerate bosons

2.2.1 Bose condensation

Noninteracting case

In this thesis, we are concerned only with degenerate bosons. In most regimes, the behaviour of identical-particle ensembles is very similar to that of distinguishable particles, and noticeable differences occur only as the system nears *degeneracy*, which is to say, when (some) single-particle modes have occupation of order $\gtrsim 1$. In such conditions, the statistics of the particles become very important. In the case of identical bosons, a statistical *clustering* effect is exhibited by the particles. The origin of this effect can be seen already in a simple toy model [164]: the distribution of two particles over two equivalent (e.g., equal energy) states. In the case of distinguishable particles, there are four ways to distinguish the particles (two choices of state for each of the two particles), and so picking permissible system states at random we expect to find both particles in the same state 50% of the time. By contrast, in the case of identical bosons, there are only three possible two-body states (both particles in state A, both in state B, and one particle in each of A and B). The probability of both particles being found in the same state is thus approximately 67% in this case. This is illustrative of a general tendency of identical bosons to exhibit an ‘attraction’ on one another, which is *purely statistical* in origin.

This behaviour is of course reflected in the equilibrium distribution of identical bosons, which is most conveniently expressed in the grand canonical ensemble, where the occupation of a single-particle state $|j\rangle$ with energy ϵ_j is

$$n(\epsilon_j) = \frac{1}{e^{\beta(\epsilon_j - \mu)} - 1}, \quad (2.10)$$

with μ the *chemical potential* and $\beta = 1/k_B T$ the inverse temperature. A crucial property of the Bose distribution (2.10) is that the excited state population $N_{\text{ex}} = \sum_{j>0} n(\epsilon_j)$

exhibits a *saturation* behaviour [47]. This can be understood by the following argument: measuring energies relative to that of the ground state (i.e. setting $\epsilon_0 = 0$), for any temperature $T > 0$ the ground state occupation becomes negative for chemical potentials $\mu > 0$. Hence we require $\mu \leq 0$. The excited state population is therefore bounded $N_{\text{ex}} \leq \sum_{k>0} [\exp(\beta\epsilon_k) - 1]^{-1} \equiv N_{\text{ex}}^{(\text{max})}$, which can be shown to be finite. Thus as atoms are added to the ensemble at fixed temperature T , eventually the excited states will admit no more atoms, and so the additional atoms accumulate in the ground state of the external potential in which they are confined. In this limit the chemical potential approaches zero from below ($\mu \rightarrow 0^-$), and the condensate occupation $n(\epsilon_0)$ diverges towards infinity. More formally, the occupation of the ground state becomes an *extensive* parameter of the system, while the occupations of all other modes remain intensive, and the ground state population is said to form a *Bose-Einstein condensate* [204, 47].

For a given number of particles N , we define the *critical temperature* T_c as that temperature at which $N_{\text{ex}}^{(\text{max})} = N$; below this temperature, atoms accumulate in the ground state, i.e., condensation occurs. A standard approach to evaluating $N_{\text{ex}}^{(\text{max})}$ is to make a *semiclassical approximation*, replacing the sum over states with an integral over energies with a density of states function¹. This approach is valid for temperatures $k_B T \gg \hbar\omega$. One obtains [204]

$$N = N_{\text{ex}}(T_c, \mu = 0) = \int_0^\infty d\epsilon \frac{g(\epsilon)}{e^{\epsilon/k_B T_c} - 1}, \quad (2.11)$$

where for the experimentally relevant case of 3D harmonic trapping, with external potential

$$V_{\text{ext}}^{(\text{harm})}(\mathbf{x}) = \frac{m}{2} (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \quad (2.12)$$

the density of states $g(\epsilon) = \epsilon^2/2(\hbar\bar{\omega})^3$, with the geometric-mean trapping frequency $\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}$. In this geometry we thus obtain the result

$$k_B T_c = \frac{\hbar\bar{\omega} N^{1/3}}{[\zeta(3)]^{1/3}} \approx 0.94 \hbar\bar{\omega} N^{1/3}, \quad (2.13)$$

where we have substituted an approximate numerical value for the *Riemann zeta* function $\zeta(z)$ evaluated at $z = 3$. We can in fact go further and calculate the number of excited particles at temperatures below T_c , where the chemical potential $\mu = 0$. We find immediately

$$N_{\text{ex}} = N \left(\frac{T}{T_c} \right)^3, \quad (2.14)$$

¹An alternative approach is to employ a series expansion for the Bose function [47], which allows one to obtain next-order corrections beyond the semiclassical limit.

and we thus have for the condensate population

$$N_0 = N \left[1 - \left(\frac{T}{T_c} \right)^3 \right]. \quad (2.15)$$

Interacting case

As the ‘clustering’ behaviour of bosons is an intrinsic quantum-statistical effect, independent of details of the atomic interaction and any external trapping potential, we expect similar condensation behaviour to occur in an interacting Bose gas. While in the non-interacting case, energy eigenstates of the many-body system are obtained as (properly symmetrised) products of single-particle eigenstates, interactions between the atoms lead to correlations in the many-body wave function, so that the concept of single-particle energy levels is not meaningful for the interacting system. In order to generalise the notion of an extensive occupation of the lowest-energy single-particle state to the interacting case, Penrose and Onsager [201] therefore considered the *one-body reduced density operator*

$$\hat{\rho}_1 = N \text{Tr}_{2\dots N} \{ \hat{\rho}_N \}, \quad (2.16)$$

which we have written here as a partial trace of the N -body density matrix $\hat{\rho}_N$, which is in general a statistical mixture $\hat{\rho}_N = \sum_s p_s |\psi_N^s\rangle \langle \psi_N^s|$ of distinct N -particle states $|\psi_N^s\rangle$. It is convenient to consider the one-body density operator in a coordinate-space representation

$$\rho_1(\mathbf{x}, \mathbf{x}') = \langle \mathbf{x}' | \hat{\rho}_1 | \mathbf{x} \rangle \equiv \langle \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}(\mathbf{x}') \rangle. \quad (2.17)$$

The one-body density matrix encapsulates all single-particle correlations in the system. For example, we can see from the definition of the extended general one-body operator (2.8) and using the cyclic property of the trace, that the expectation of any extended one-body operator \hat{J} is simply given by the trace of the original single-body operator with respect to the one-body density matrix, i.e. $\langle \hat{J} \rangle = \text{Tr} \{ \rho_1 J \}$. As the one-body density matrix is Hermitian, it can be diagonalised with *real* eigenvalues n_i , and corresponding eigenvectors $\chi_i(\mathbf{x})$

$$\rho_1(\mathbf{x}, \mathbf{x}') = \sum_i n_i \chi_i^*(\mathbf{x}) \chi_i(\mathbf{x}'). \quad (2.18)$$

It is immediately apparent that the eigenvalue n_i quantifies the (mean) occupation of the mode $\chi_i(\mathbf{x})$ in the many-body state. An appropriate generalisation of the notion of Bose-Einstein condensation is thus obtained by regarding the most highly occupied mode $\chi_0(\mathbf{x})$ as the appropriate analogue of the lowest-energy single-particle state of the noninteracting system: when the occupation n_0 is an extensive parameter, with all others

remaining intensive, a condensate is formed, with *condensate orbital* $\chi_0(\mathbf{x})$. In this case we can separate out the condensate and write

$$\rho_1(\mathbf{x}, \mathbf{x}') = n_0 \chi_0^*(\mathbf{x}) \chi_0(\mathbf{x}') + \sum_{i>0} n_i \chi_i^*(\mathbf{x}) \chi_i(\mathbf{x}'). \quad (2.19)$$

In the thermodynamic limit we can replace the sum by an integral, which tends to zero for large separations $|\mathbf{x} - \mathbf{x}'|$ [207]. By contrast, the contribution of the condensate remains finite, provided \mathbf{x} and \mathbf{x}' remain within the extent of the condensate orbital, and in the limiting case of an infinite, homogeneous system we obtain the result that $\lim_{|\mathbf{x}-\mathbf{x}'|\rightarrow\infty} \rho_1(\mathbf{x}, \mathbf{x}') \rightarrow n_0$, a property known as *off-diagonal long-range order* [201]. In quantum optical terms, this property reflects the first-order coherence of the Bose field, as discussed by Naraschewski and Glauber [186].

It should be noted, however, that although this Penrose-Onsager definition of condensation is unambiguous in many simple scenarios, it is not difficult to construct systems for which the one-body density matrix can give misleading results. Examples include strongly nonequilibrium systems [164] and situations involving attractive interactions [265], and indeed even the (common) case of harmonic trapping allows undamped centre-of-mass motion of the atomic cloud [72]², which can lead to questionable results for the condensate fraction from the Penrose-Onsager approach [76]. The appropriate definition of condensation in more general situations remains an uncertain and controversial topic, as discussed in references [164, 265, 76, 203, 98, 76, 271]. The characterisation of coherence and condensation in the atomic field beyond the simple Penrose-Onsager definition forms a major theme of this thesis, and we discuss alternative measures of coherence in chapters 5, 6, 7, and 9.

2.2.2 The Gross-Pitaevskii equation

Having determined the appropriate definition of the condensate in terms of the condensate orbital $\chi_0(\mathbf{x})$, we wish to *characterise* the orbital $\chi_0(\mathbf{x})$ and its evolution with time. At zero temperature, the form and evolution of $\chi_0(\mathbf{x})$ are governed by the *Gross-Pitaevskii* equation (GPE), which we derive here. In the conditions of condensation in this zero-temperature limit, only low-energy, binary interactions between atoms are relevant [60], and these interactions are well characterised by a single parameter, the *s*-wave scatter-

²A very general result for harmonically confined interacting-particle systems guarantees the presence of dipole oscillations of the centre-of-mass of the particle assembly at the appropriate trapping frequencies [72]. This is in close analogy with *Kohn theorem* for the centre-of-mass cyclotron resonance of interacting electrons in a static magnetic field [149], and it is thus conventional to speak of the ‘Kohn theorem’ and ‘Kohn modes’ of trapped degenerate gases.

ing length, independent of the details of the interatomic potential. This motivates the replacement of the interatomic potential by a purely repulsive *effective potential* [207] or a (regularised) delta-function *pseudopotential* [132] which produces the same s -wave scattering length. Quite aside from the fact that the precise details of the interatomic potential are unknown, this replacement is essential as the true interatomic potential supports bound states, precluding the perturbative treatment of interactions which is implicit in mean-field treatments [47]. For the purposes of our discussion here it will be sufficient to assume the common case in which the inter-particle interactions are described by the simple contact interaction $V(\mathbf{x}) = U_0\delta(\mathbf{x})$, where the interaction coefficient $U_0 \equiv 4\pi\hbar^2 a_s/m$, with a_s the s -wave scattering length. Substituting this effective potential into the second-quantised Hamiltonian (2.6), we obtain the cold-collision Hamiltonian

$$\hat{H} = \int d\mathbf{x} \hat{\Psi}^\dagger(\mathbf{x}) H_{\text{sp}} \hat{\Psi}(\mathbf{x}) + \frac{U_0}{2} \int d\mathbf{x} \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}(\mathbf{x}) \hat{\Psi}(\mathbf{x}). \quad (2.20)$$

Hartree-Fock ansatz

We first consider the time-independent GPE, which describes the form of the stationary condensed mode $\chi_0(\mathbf{x})$ as a nonlinear eigenvalue problem. There are several ways to derive this equation; to best give insight into the physical content of the equation, we derive it here by assuming a Hartree-Fock ansatz for the many-body wave function in terms of the condensate orbital $\chi_0(\mathbf{x})$. Specifically, we assume the ansatz [164]

$$\psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \prod_{i=1}^N \chi_0(\mathbf{x}_i). \quad (2.21)$$

This ansatz is clearly in accord with our expectations for a zero-temperature condensate: every atom resides in the same single-particle orbital $\chi_0(\mathbf{x})$. The energy of the system (i.e., the expectation value of (2.20) in the corresponding Fock state) is

$$E_{\text{HF}} = N \int d\mathbf{x} \left[\frac{\hbar^2}{2m} |\nabla \chi_0(\mathbf{x})|^2 + V_{\text{ext}}(\mathbf{x}) |\chi_0(\mathbf{x})|^2 + \frac{(N-1)U_0}{2} |\chi_0(\mathbf{x})|^4 \right]. \quad (2.22)$$

By functionally minimising this energy (see, e.g., [26]), subject to the constraint of normalisation of the condensate orbital ($\int d\mathbf{x} |\chi_0(\mathbf{x})|^2 = 1$), we obtain the time-independent GPE

$$-\frac{\hbar^2}{2m} \nabla^2 \chi_0(\mathbf{x}) + V_{\text{ext}}(\mathbf{x}) \chi_0(\mathbf{x}) + (N-1)U_0 |\chi_0(\mathbf{x})|^2 \chi_0(\mathbf{x}) = \lambda \chi_0(\mathbf{x}), \quad (2.23)$$

where the parameter λ here enters as the Lagrange multiplier associated with the normalisation constraint [164, 26], and it can be shown that $\lambda = \partial E_{\text{HF}}/\partial N$, i.e., λ is the

chemical potential of the condensate. The term in $|\chi_0(\mathbf{x})|^2$ describes the *mean-field potential* experienced by each atom due to the presence of all the other atoms, and we see that it vanishes (properly) in the case of a single atom. In practice we typically approximate $N - 1 \approx N$ on the basis that N is large, and the GPE is traditionally written in terms of the function $\phi(\mathbf{x}) \equiv \sqrt{N}\chi_0(\mathbf{x})$ with norm square equal to the total particle number [164]. The function $\phi(\mathbf{x})$ is also commonly referred to as the condensate ‘wave function’, though one should be careful not to confuse it with the true many-body wave function (2.21) of the condensate in this approach. Although this Hartree-Fock approach yields the clearest physical insight into the Gross-Pitaevskii (GP) equation, the corresponding time-dependent GP equation cannot be derived rigorously in this manner [164], and so we must consider less transparent approaches to defining the condensate wave function.

Symmetry-breaking approach

A widespread approach to deriving the Gross-Pitaevskii equation is based on the idea of *spontaneous symmetry breaking* by a phase transition, which although less transparent than the Hartree-Fock approach, is more direct. Formally, such a phase transition is described in terms of the emergence of a nonzero *order parameter* below the phase transition. The archetypal example of such an order parameter is the magnetisation $\langle \mathbf{M} \rangle$ which emerges spontaneously in a ferromagnetic system cooled below the Curie temperature [157]. Conventionally, Bose-Einstein condensation is formulated as a similar symmetry-breaking phase transition, in which the emergent order parameter is the expectation value of the field operator $\langle \hat{\Psi}(\mathbf{x}) \rangle$, which is associated with condensation in the system. As the first moment $\langle \hat{\Psi}(\mathbf{x}) \rangle$ of the field acquires a finite value, its phase becomes well-defined, breaking the U(1) (gauge) symmetry of the Hamiltonian (2.20) under global rotations of the field operator phase $\hat{\Psi}(\mathbf{x}) \rightarrow \hat{\Psi}(\mathbf{x})e^{i\theta}$ (see, e.g., [161]). It is important to note that rotations of this phase are generated by the *total particle number* $\hat{N} = \int d\mathbf{x} \hat{\psi}^\dagger(\mathbf{x})\hat{\psi}(\mathbf{x})$ [25], and the breaking of this symmetry is thus tantamount to allowing the nonconservation of particle number. This symmetry breaking is therefore *not physical*, and is merely a computational convenience. Formally this fiction is introduced by calculating expectation values in a modified ensemble [131] of states with only a small dispersion in the phase (and amplitude) of the order parameter. We can achieve a similar result at zero temperature by assuming that the atomic field is in a *coherent state* $|\phi(\mathbf{x})\rangle$ [106], which is a superposition of states with *all* (physical) particle numbers, such that $\hat{\Psi}(\mathbf{x})|\phi(\mathbf{x})\rangle = \phi(\mathbf{x})|\phi(\mathbf{x})\rangle$.

The time-dependent Gross-Pitaevskii equation is derived using this notion of sponta-

Chapter 2. Background theory

neous symmetry breaking as follows: using the commutation relation (2.4) we derive from the second-quantised Hamiltonian (2.20) the Heisenberg equation of motion

$$i\hbar \frac{\partial \hat{\Psi}(\mathbf{x})}{\partial t} = [\hat{\Psi}(\mathbf{x}), \hat{H}] = \{H_{\text{sp}} + U_0 \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}(\mathbf{x})\} \hat{\Psi}(\mathbf{x}). \quad (2.24)$$

Taking the expectation value of (2.24) in the coherent state $|\phi(\mathbf{x})\rangle$ amounts to making the replacement $\hat{\Psi}(\mathbf{x}) \rightarrow \phi(\mathbf{x})$, yielding the time-dependent Gross-Pitaevskii equation

$$i\hbar \frac{\partial \phi(\mathbf{x}, t)}{\partial t} = \{H_{\text{sp}} + U_0 |\phi(\mathbf{x}, t)|^2\} \phi(\mathbf{x}, t). \quad (2.25)$$

This equation describes the *collective* motion of the condensate at zero temperature, as for example in response to a time-dependent external trapping potential $V_{\text{ext}}(\mathbf{x}, t)$. It is important to note that this equation only describes the motion so long as the Bose field remains completely condensed. However, *dynamical instabilities* of the condensate can arise, and are marked in the GP solution by the instability of the condensate mode to exponential growth of collective excitations. This in fact represents the runaway amplification of beyond-GP fluctuations of the Bose field, and so the appearance of a dynamical instability marks the breakdown of the Gross-Pitaevskii description, as we discuss further in section 2.2.3.

The derivation of the time-independent GP equation in the symmetry-breaking approach parallels the Hartree-Fock derivation of section 2.2.2, as we similarly obtain the GP equation by minimising the energy of the many-body system subject to an ansatz for the many-body state. As we have broken the conservation of particle number, we work in the grand-canonical ensemble where the *mean* atom number is conserved. We thus consider the grand-canonical Hamiltonian

$$\hat{K} = \hat{H} - \mu \hat{N}, \quad (2.26)$$

where \hat{H} is the Hamiltonian (2.20), and the chemical potential μ enforces the conservation of particle number on average. Taking the expectation of equation (2.26) in the state $|\phi(\mathbf{x})\rangle$ yields [86]

$$K_0 = H_0 - \mu N_0 = \int d\mathbf{x} \left[\phi^*(\mathbf{x}) (H_{\text{sp}} - \mu) \phi(\mathbf{x}) + \frac{U_0}{2} |\phi(\mathbf{x})|^4 \right]. \quad (2.27)$$

We will refer to the functional $H_0[\phi(\mathbf{x})]$ as the Gross-Pitaevskii *energy functional*. In the zero-temperature limit appropriate to our assumed coherent state $|\phi(\mathbf{x})\rangle$, the free energy of the system is simply $\langle \hat{K} \rangle$, and so the system equilibrium is found by minimising

equation (2.27), i.e., by requiring that K_0 is stationary under variations of $\phi(\mathbf{x})$ and its complex conjugate. The Euler-Lagrange equation resulting from variations of the latter is the time-independent GP equation

$$\{H_{\text{sp}} + U_0|\phi(\mathbf{x})|^2\} \phi(\mathbf{x}) = \mu\phi(\mathbf{x}). \quad (2.28)$$

We note that this form of the time-independent GPE can also be obtained from the time-dependent GPE (2.25) by assuming a uniform phase rotation $\phi(\mathbf{x}, t) = e^{-i\mu t}\phi(\mathbf{x})$. We notice that whereas the time development of a true quantum-mechanical wave function is governed by its energy, of the phase of the order parameter rotates with frequency determined by its *chemical potential* [207].

Healing length

An important length scale which characterises solutions of the GP equation is the *healing length*, which is the distance over which the condensate wave function, constrained to zero density by some applied potential, returns to its bulk density. This length is thus set by the balance of the kinetic energy associated with the curvature of the wave function and the interaction energy associated with its density [204]. Equating the two we find the healing length

$$\eta = \frac{\hbar}{\sqrt{2mn(\mathbf{x})U_0}}, \quad (2.29)$$

where $n(\mathbf{x}) = |\phi(\mathbf{x})|^2$ is the local density of the condensate. In inhomogeneous condensates the healing length is thus a spatially varying quantity. However, in characterising inhomogeneous condensates it is common to simply use the healing length corresponding to the central (i.e., peak) density of the condensate as an order-of-magnitude estimate [60]. This healing length characterises (for example) the core size of a *quantum vortex* in the condensate (section 2.3.2).

Thomas-Fermi limit

In the limit of large condensate population, the nonlinearity in the time-independent Schrödinger equation dominates, and the kinetic energy becomes negligible in comparison. In this limit we may therefore simply drop the curvature term from equation (2.28), leaving us with an algebraic equation for the condensate density, which we can rearrange to find the *Thomas-Fermi condensate wave function*

$$\Psi_{\text{TF}}(\mathbf{x}) = \sqrt{\frac{\mu_{\text{TF}} - V_{\text{ext}}(\mathbf{x})}{U_0}} \Theta(\mu_{\text{TF}} - V_{\text{ext}}(\mathbf{x})), \quad (2.30)$$

Chapter 2. Background theory

where $\Theta(x)$ denotes a Heaviside function. The quantity μ_{TF} is the Thomas-Fermi (TF) limit chemical potential, which for harmonic trapping becomes [189]

$$\mu_{\text{TF}} = \frac{1}{2} (15a\hbar^2 \sqrt{m\omega_x\omega_y\omega_z} N_0)^{2/5}, \quad (2.31)$$

where $N_0 = \int d\mathbf{x} |\Psi_{\text{TF}}|^2$ is the Thomas-Fermi condensate population. In this geometry, the half-width (Thomas-Fermi radius) of the TF condensate orbital in the i^{th} radial direction is

$$R_i^{\text{TF}} = \frac{1}{\omega_i} \sqrt{\frac{2\mu_{\text{TF}}}{m}}. \quad (2.32)$$

For large condensates the Thomas-Fermi wave function provides an excellent approximation to the exact Gross-Pitaevskii condensate mode, differing only at the condensate boundary, where the density $|\Psi_{\text{TF}}(\mathbf{x})|^2$ of the TF condensate vanishes abruptly at $x_i = R_i^{\text{TF}}$, whereas the boundary of the true Gross-Pitaevskii condensate orbital is smoothed so as to reduce its kinetic energy.

Hydrodynamic formulation and collective modes

We now show that the time-dependent Gross-Pitaevskii equation (2.25) can be reformulated as a hydrodynamic equation for the condensate motion [249, 60]. We perform the *Madelung transformation*, making the substitution $\phi(\mathbf{x}) = \sqrt{n(\mathbf{x})} e^{iS(\mathbf{x})}$ for the condensate wave function. In terms of the hydrodynamic velocity $\mathbf{v} = (\hbar/m)\nabla S$ we thus obtain the *continuity equation*

$$\frac{\partial}{\partial t} n + \nabla \cdot (n\mathbf{v}) = 0, \quad (2.33)$$

expressing the conservation of the condensate population, and the *Euler equation*

$$m \frac{\partial}{\partial t} \mathbf{v} + \nabla \left(V_{\text{ext}} + gn - \frac{\hbar^2}{2m\sqrt{n}} \nabla^2 \sqrt{n} + \frac{mv^2}{2} \right) = 0. \quad (2.34)$$

The term in $\nabla^2 \sqrt{n}$ here is the *quantum pressure* resulting from the kinetic energy term in the time-dependent Gross-Pitaevskii equation. In the limit of large condensates we can drop the quantum pressure term, which corresponds to the Thomas-Fermi approximation discussed above, to obtain the simpler Euler equation

$$m \frac{\partial}{\partial t} \mathbf{v} + \nabla \left(V_{\text{ext}} + gn + \frac{mv^2}{2} \right) = 0. \quad (2.35)$$

One can easily show that the time-independent solution of equation (2.35) corresponds to the Thomas-Fermi solution discussed previously [60]. Introducing a small deviation

$\delta n(\mathbf{x}, t) = n(\mathbf{x}, t) - n_{\text{TF}}(\mathbf{x})$, and linearising equations (2.33) and (2.35) about the TF solution, we can estimate the frequencies of collective modes of oscillation of the condensate [249]. In the case of harmonic trapping axially symmetric along the z axis (most relevant to this thesis), we obtain an equation

$$m \frac{\partial^2}{\partial t^2} \delta n = \nabla \cdot \left\{ \left[\mu - \frac{m}{2} (\omega_{\perp}^2 r_{\perp}^2 + \omega_z^2 z^2) \right] \nabla \delta n \right\}, \quad (2.36)$$

for the linear deviations. Substituting a time-harmonic ansatz $\delta n(\mathbf{x}, t) = \delta n(\mathbf{x}) e^{-i\omega t}$ into equation (2.36), we can obtain a differential equation for the spatial modes of oscillation, and deduce the corresponding oscillation frequencies [249]. Of particular interest are the so-called *surface modes*, which have no radial nodes. In the axially symmetric geometry, each such mode is characterised by the quantum numbers l and m which specify the mode angular momentum and its projection along the z axis, respectively. In particular, we find frequencies $\omega_{l,m}$ for the *dipole oscillations*

$$\omega_{1,1} = \omega_{\perp}; \quad \omega_{1,0} = \omega_z, \quad (2.37)$$

which are of course the *exact* eigenfrequencies of the dipole modes of the harmonically trapped interacting gas by the Kohn theorem (section 2.2.1). We also find in general

$$\omega_{l,\pm l} = \sqrt{l} \omega_{\perp}, \quad (2.38)$$

which is the frequency of a surface oscillation of multipolarity l about the z axis. Such a surface oscillation can be excited by a rotating trap anisotropy of the same multipolarity. In general, the efficiency of this excitation frequency will depend on the frequency at which the anisotropy is rotated. Taking into account the multipolarity of the trap we expect *resonant* excitation to occur at a frequency $\omega_{\text{res}} = \omega_l/l$. In the particular case of quadrupole mode excitation relevant to chapters 8 and 9 we find the resonant driving frequency

$$\omega_Q = \omega_{\perp}/\sqrt{2}. \quad (2.39)$$

2.2.3 Bogoliubov theory

In this section we consider the character of excitations of the condensed field. We therefore assume that a single mode ($\chi_0(\mathbf{x})$) of the atomic field is condensed, and replace the corresponding annihilation operator $\hat{a}_0 = \int d\mathbf{x} \chi_0^*(\mathbf{x}) \hat{\Psi}(\mathbf{x})$ with the *number* $\sqrt{N_0}$ (where N_0 is the condensate occupation), and retain a second-quantised description for the remainder of the field. Although this is the standard symmetry-breaking approach in the literature,

we note that Castin and Dum [50] have provided a more rigorous and less conceptually ambiguous *number-conserving* formulation of the Bogoliubov expansion (see also [100]). However, the additional complexity involved in the rigorous asymptotic expansion of [50] obscures somewhat the simple physical picture of the Bogoliubov expansion we wish to present here. We thus make the *Bogoliubov shift*

$$\hat{\Psi}(\mathbf{x}) \rightarrow \phi(\mathbf{x}) + \hat{\delta}(\mathbf{x}), \quad (2.40)$$

separating the field operator $\hat{\Psi}(\mathbf{x})$ into a condensed part $\phi(\mathbf{x}) = \sqrt{N_0}\chi_0(\mathbf{x})$ and a *fluctuation operator* $\hat{\delta}(\mathbf{x})$, which is assumed to obey the standard bosonic commutation relations (equation (2.4)). We then proceed to derive a Hamiltonian which describes the fluctuation operator $\hat{\delta}(\mathbf{x})$, and thus characterises the excitations of the condensate. Substituting the separation (2.40) into the second-quantised Hamiltonian (2.20) and assuming that the condensate wave function $\phi(\mathbf{x})$ is a solution of (2.23), we obtain

$$\begin{aligned} \hat{K} = & \int d\mathbf{x} \hat{\delta}^\dagger \left[H_{\text{sp}} + 2U_0|\phi|^2 - \mu \right] \hat{\delta} \\ & + \frac{U_0}{2} \int d\mathbf{x} \left[(\phi^*)^2 \hat{\delta}\hat{\delta} + \phi^2 \hat{\delta}^\dagger \hat{\delta}^\dagger + \phi^* \hat{\delta}^\dagger \hat{\delta}\hat{\delta} + \phi \hat{\delta} \hat{\delta}^\dagger \hat{\delta} + \hat{\delta}^\dagger \hat{\delta} \hat{\delta}\hat{\delta} \right], \end{aligned} \quad (2.41)$$

where we have dropped a constant term K_0 which is simply an energy shift and does not otherwise affect the motion of the fluctuation operator³, and we suppress the spatial arguments of ϕ and $\hat{\delta}$ for notational convenience. In general, this Hamiltonian remains, like the original second-quantised Hamiltonian (2.20), intractable. However, a Hamiltonian which is quadratic in the field operators can always be diagonalised (in a generalised sense, as we discuss below) [25] and so to proceed from here we approximate this Hamiltonian by one in which higher powers of $\hat{\delta}$ do not occur. The approach we take is to simply neglect terms which are cubic and quartic in $\hat{\delta}$. This approximation is clearly appropriate in the limit that the noncondensed fraction of the field is small compared to the condensate (i.e. $\langle \hat{\delta}^\dagger \hat{\delta} \rangle \ll |\phi|^2$), and we obtain in this case the quadratic Hamiltonian

$$\hat{K}_{\text{quad}} = \int d\mathbf{x} \hat{\delta}^\dagger \left[H_{\text{sp}} + 2U_0|\phi|^2 - \mu \right] \hat{\delta} + \frac{U_0}{2} \int d\mathbf{x} \left[(\phi^*)^2 \hat{\delta}\hat{\delta} + \phi^2 \hat{\delta}^\dagger \hat{\delta}^\dagger \right]. \quad (2.42)$$

As this Hamiltonian contains terms like $\hat{\delta}\hat{\delta}$, its eigenbasis is *not* a single-particle basis, but can be expressed in terms of two-component (*spinor*) basis functions. Physically, this is because (at this order) the excitations of the condensate involve *pairs* of atoms: the interactions ‘mix’ the roles of creation and annihilation operators of single-particle

³Linear terms in $\hat{\delta}$ vanish due to the stationarity of (2.27) under variations of $\phi(\mathbf{x})$ [86].

modes, and the elementary excitations of the system are created by so-called *quasiparticle operators*, which are superpositions of these creation and annihilation operators [25]. Such quasiparticle operators are introduced by defining

$$\hat{b}_k = \int d\mathbf{x} \left[u_k^*(\mathbf{x}) \hat{\delta}(\mathbf{x}) - v_k^*(\mathbf{x}) \hat{\delta}^\dagger(\mathbf{x}) \right]. \quad (2.43)$$

Requiring that the \hat{b}_k satisfy the standard Bose commutation relations ($[\hat{b}_j, \hat{b}_k^\dagger] = \delta_{jk}$; $[\hat{b}_j, \hat{b}_k] = 0$) we find that the quasiparticle spinors $(u_k(\mathbf{x}), v_k(\mathbf{x}))^T$ are orthogonal with respect to the metric

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.44)$$

i.e., we have

$$\int d\mathbf{x} (u_j^*(\mathbf{x}), v_j^*(\mathbf{x})) \cdot \sigma_z \cdot \begin{pmatrix} u_k(\mathbf{x}) \\ v_k(\mathbf{x}) \end{pmatrix} = \int d\mathbf{x} u_j^* u_k - v_j^* v_k = \delta_{jk}. \quad (2.45)$$

Our aim is to find a *diagonal Hamiltonian*, written in terms of the \hat{b}_k and their Hermitian conjugates, i.e., one which takes the form (up to a constant energy)

$$\hat{K}_{\text{quad}} = \sum_k \epsilon_k \hat{b}_k^\dagger \hat{b}_k. \quad (2.46)$$

This form for the Hamiltonian implies the commutation relation $[\hat{b}_k, \hat{K}_{\text{quad}}] = \epsilon_k \hat{b}_k$. Evaluating this commutator explicitly for the Hamiltonian (2.42), we find that the mode functions $u_k(\mathbf{x}), v_k(\mathbf{x})$ satisfy the *Bogoliubov-de Gennes* equations

$$\mathcal{L} \begin{pmatrix} u_k(\mathbf{x}) \\ v_k(\mathbf{x}) \end{pmatrix} = \epsilon_k \begin{pmatrix} u_k(\mathbf{x}) \\ v_k(\mathbf{x}) \end{pmatrix}, \quad (2.47)$$

which we have written here in terms of the matrix

$$\mathcal{L} \equiv \begin{pmatrix} H_{\text{sp}} - \mu + 2U_0|\phi|^2 & U_0\phi^2 \\ -U_0(\phi^*)^2 & -[H_{\text{sp}} - \mu + 2U_0|\phi|^2]^* \end{pmatrix}. \quad (2.48)$$

Dual modes and Goldstone mode

A symmetry property of the matrix (2.48) implies that for every spinor $(u_k(\mathbf{x}), v_k(\mathbf{x}))^T$ which is a solution of equation (2.47) (with energy ϵ_k), there is a dual solution $(v_k^*(\mathbf{x}), u_k^*(\mathbf{x}))^T$ with energy $-\epsilon_k^*$, and which clearly has normalisation (with respect to the inner product (2.45)) -1 [47]. These negative-normed modes are simply a dual rep-

resentation of the positive-normed modes, and do not appear in the Bogoliubov form of the quadratic Hamiltonian. However, these two families of modes do not in fact exhaust the eigenbasis of \mathcal{L} ; the spinor

$$\Theta_G \equiv \begin{pmatrix} \phi(\mathbf{x}) \\ -\phi^*(\mathbf{x}) \end{pmatrix}, \quad (2.49)$$

is an eigenvector of \mathcal{L} , with eigenvalue $\epsilon_G = 0$, and normalisation $\int d\mathbf{x} \Theta_G^\dagger \sigma_z \Theta_G = 0$. This mode does not appear in the diagonalised Hamiltonian (2.46) due to its zero eigenvalue, and is the *spurious* or *Goldstone* mode associated with the breaking of phase symmetry [25]. A more careful treatment of the problem [165] reveals that the full quadratic Hamiltonian is of the form

$$H_{\text{quad}} = \sum_k \epsilon_k \hat{b}_k^\dagger \hat{b}_k + \frac{\hat{P}^2}{2m_G}, \quad (2.50)$$

where the ‘momentum’ operator $\hat{P} = \int d\mathbf{x} [\chi_0^*(\mathbf{x}) \hat{\delta}(\mathbf{x}) + \chi_0(\mathbf{x}) \hat{\delta}^\dagger(\mathbf{x})]$ is related to the phase of the condensate mode, and the ‘mass’ m_G is inversely proportional to the peak density of the condensate [165]. The appearance here of a momentum variable without a corresponding potential signifies that the phase of the condensate diffuses over time, restoring the symmetry spuriously broken by our coherent-state approximation [165]. In general, such a term appears for each symmetry of the Hamiltonian which is broken by the assumed vacuum state of the approximate quadratic Hamiltonian [25]. It is important to note, however, that this particular spurious mode arises due to the breaking of phase symmetry, which, as we have noted, is not physical. Indeed, no such phase-diffusion term arises in the more careful number conserving approach of [50]. In fact, for finite systems such as we discuss in this thesis, broken symmetries arise only as a result of approximations [25]. The form of the Goldstone Hamiltonian term $\hat{P}^2/2m_G$ appearing here as a result of phase-symmetry breaking will allow us to better understand the consequences of phase and rotational symmetry breaking in our classical-field simulations in chapters 5 and 6 respectively.

Relation to small oscillations in Gross-Pitaevskii theory

We now illustrate a deep connection between the Bogoliubov theory for the elementary excitations of the condensed Bose gas in second quantisation, and the small oscillations (‘normal modes’) of the condensate in the time-dependent Gross-Pitaevskii formalism. We consider the small oscillations of the condensate in the Gross-Pitaevskii approach, generalising the hydrodynamic approach to the oscillations previously discussed in section 2.2.2. We assume some ‘reference’ configuration $\phi_0(\mathbf{x}, t)$ of the time-dependent Gross-

Pitaevskii equation (2.25), and calculate to linear order the evolution of a small deviation $\delta\phi(\mathbf{x}, t) = \phi(\mathbf{x}, t) - \phi_0(\mathbf{x}, t)$ about this reference solution. We thus obtain the linear evolution equation [47]

$$i\hbar \frac{\partial}{\partial t} \delta\phi(\mathbf{x}, t) = -\frac{\hbar^2}{2m} \nabla^2 \delta\phi(\mathbf{x}, t) + V_{\text{ext}}(\mathbf{x}) \delta\phi(\mathbf{x}, t) + 2U_0 |\phi_0(\mathbf{x}, t)|^2 \delta\phi(\mathbf{x}, t) + U_0 \phi_0^2(\mathbf{x}, t) \delta\phi^*(\mathbf{x}, t). \quad (2.51)$$

We consider here the particular case in which $\phi_0(\mathbf{x}, t)$ corresponds to an *eigenfunction* of the time-independent Gross-Pitaevskii equation, i.e., $\phi_0(\mathbf{x}, t) = \phi_0(\mathbf{x}) e^{-i\mu t/\hbar}$, and assume the ansatz

$$\delta\phi(\mathbf{x}, t) = e^{-i\mu t/\hbar} \sum_k [u_k(\mathbf{x}) e^{-i\omega_k t} + v_k^*(\mathbf{x}) e^{i\omega_k t}]. \quad (2.52)$$

Substituting the ansatz (2.52) into equation (2.51), and equating powers of $e^{\pm i\omega_k t}$, we find that the mode functions $u_k(\mathbf{x}), v_k(\mathbf{x})$ are solutions of

$$\mathcal{L} \begin{pmatrix} u_k(\mathbf{x}) \\ v_k(\mathbf{x}) \end{pmatrix} = \hbar\omega_k \begin{pmatrix} u_k(\mathbf{x}) \\ v_k(\mathbf{x}) \end{pmatrix}, \quad (2.53)$$

where \mathcal{L} is given by equation (2.48). We thus have the interesting result that the small *collective* oscillations of the condensate, as described by Gross-Pitaevskii equation, are in one-to-one correspondence with the *elementary excitations* of the system in the quadratic (Bogoliubov) approximation to the second-quantised field theory, with the frequencies ω_k of the small oscillations equal to the (Heisenberg-picture) phase rotation frequencies ϵ_k/\hbar of the Bogoliubov ladder operators. This is in some sense not surprising, as the above linearisation of the Gross-Pitaevskii equation amounts to a quadratisation of the Gross-Pitaevskii *energy functional* (2.27) about the Gross-Pitaevskii eigenfunction solution, in direct analogy to the quadratisation of the second-quantised Hamiltonian we performed to obtain the elementary excitations.

Stability of excitations

In equilibrium, we think of the condensate as being the ‘ground’ (or *vacuum*) state of the system, with the weakly populated Bogoliubov modes constituting some excitation of the system above this ground state, due to thermal (and quantum) fluctuations. In such a situation, the Bogoliubov excitation energies ϵ_k are therefore real and positive quantities. However, in more general situations, the excitation energies may fail to be positive or even real, and this signifies an instability of the condensate, as we now discuss.

Thermodynamic instability

It is possible that the Gross-Pitaevskii wave function $\phi(\mathbf{x})$, although an eigenfunction of the Gross-Pitaevskii equation, is not the lowest-energy eigenfunction (for its occupation N), i.e., it is a collectively excited state of the condensate. It is then possible that one or more Bogoliubov energies ϵ_k may be negative. In the presence of dissipation, the condensate can then lose energy as the population of the negative-energy Bogoliubov mode increases, and this can lead to a complete restructuring of the degenerate gas. In this case the condensate is considered to be *thermodynamically unstable*. It should be noted however that a GP mode can be an excited eigenfunction but nevertheless possess no negative-energy Bogoliubov excitations, i.e., if it resides at a *local* minimum of the GP energy functional (2.27). In this case the state is referred to as *thermodynamically metastable*.

Dynamic instability

A second case is that in which an eigenvalue ϵ_k has a nonzero imaginary part. Clearly an energy eigenvalue with a positive imaginary part yields exponential growth of the mode population with time, and it can be shown [47] that if ϵ_k is an eigenvector of \mathcal{L} then ϵ_k^* is also, and we thus require that there are no eigenvalues with negative imaginary parts either. In the case that such a complex eigenvalue appears in the Bogoliubov spectrum, the condensate mode is said to be *dynamically unstable*. It is important to note that the growth of population in a dynamically unstable mode does *not* require dissipation, but is rather a dynamical amplification of population in the unstable mode, fed by the condensate. In the second-quantised picture, dynamically unstable modes will always undergo growth due to their irreducible *vacuum* population. By the correspondence between the Bogoliubov modes and the small oscillations of the condensate, the corresponding oscillations of the Gross-Pitaevskii equation will also undergo unstable growth with time, provided that they are populated by (e.g.) numerical noise. The evolution in this case can quickly lead to field dynamics beyond the validity not only of the GP equation but also of the Bogoliubov description. It is intuitive that a dynamical instability can only occur if the condensate mode is not the ground Gross-Pitaevskii eigenstate, and indeed it can easily be shown that thermodynamic stability implies dynamic stability [47]. In fact, it can be shown that dynamical instabilities arise from a kind of generalised avoided level crossing occurring between excitations with positive and negative energies [241, 12].

Self-consistent mean-field theories

In order to describe condensates at higher temperatures in such a mean-field framework, we must (by some suitable approximation) retain the effects of terms in the grand canonical Hamiltonian (2.41) of cubic and quadratic order in $\hat{\delta}(\mathbf{x})$. One approach is to assume essentially *ad hoc* factorisations of the field operators, such as $\hat{\delta}^\dagger \hat{\delta} \hat{\delta} \rightarrow 2\langle \hat{\delta}^\dagger \hat{\delta} \rangle \hat{\delta} + \langle \hat{\delta} \hat{\delta} \rangle \hat{\delta}^\dagger$ [119]. The second expectation value in this approximation is the *anomalous density* $m(\mathbf{x}) \equiv \langle \hat{\delta}(\mathbf{x}) \hat{\delta}(\mathbf{x}) \rangle$ which represents the pairing correlations in the noncondensed component induced by the interacting condensate. Using such Hartree-Fock-Bogoliubov (HFB) approximations, one obtains *generalised* Gross-Pitaevskii and Bogoliubov-de Gennes equations, which include additional mean-field potential terms depending on the noncondensed population and anomalous density, and which must therefore be solved *self-consistently* (i.e., by an iterative procedure [25, 119]). Such approaches have been applied to finite-temperature equilibrium scenarios [113, 135, 109]. There are, however, nontrivial complications introduced by the factorisation approximations [119, 183, 217]. Moreover, these theories do not allow for any particle exchange between the condensate and thermal cloud [217], which severely limits their applicability in nonequilibrium scenarios. The most conspicuous failing of the mean-field theories, however, is the requirement that the Bose field be explicitly separated into condensed and noncondensed parts. In general nonequilibrium scenarios (such as, e.g., turbulent dynamics of the Bose field [170]) no such clear distinction can be made, and so mean-field theories are inapplicable in such regimes.

Number-conserving approaches

As we have noted, it is possible to reformulate the Bogoliubov theory in terms which explicitly preserve the U(1) symmetry of the fundamental second-quantised Hamiltonian (2.20), and thus conserve particle number [115, 50, 100, 183, 107], at the expense of a significant increase in complexity. In these approaches the condensate is regarded as the dominant eigenvector of the one-body density matrix $\chi_0(\mathbf{x})$, with corresponding annihilation operator \hat{a}_0 , and the noncondensed field is described in terms of operators which commute with the total particle number; e.g., Castin and Dum [50] work in terms of the noncondensate operator $\hat{\Lambda} = \hat{N}^{-1/2} \hat{a}_0^\dagger \hat{\delta}$, where \hat{a}_0^\dagger creates a quanta in the condensate orbital $\chi_0(\mathbf{x})$, and similar number-conserving operators are employed in references [100, 183, 107]. In particular, the approaches of references [100, 50] demonstrate that the time-dependent Gross-Pitaevskii equation traditionally derived under symmetry-breaking assumptions (section 2.2.2) is indeed valid at zero temperature. An essential aspect of these treatments is that the rigorous separation of the Bose field operator into condensate and noncondensate parts requires that the two components are *orthogonal* (i.e.,

$\int d\mathbf{x} \chi_0^*(\mathbf{x}) \hat{\delta}(\mathbf{x}) = 0$)⁴. Time-dependent formalisms in this number-conserving approach have been formulated both for low temperature [50, 100] and moderate temperature regimes [107]. The method of [107] requires only that the condensed atoms constitute the *majority* of the field population, and allows in principle for (e.g.) the growth of the thermal component during the evolution. However, the method is complicated significantly by the appearance of nonlocal terms which serve to maintain the orthogonality of the time-dependent condensed and noncondensed components of the field throughout their evolution, and has therefore only been applied in the linear-response regime so far [217]. Furthermore, such an approach suffers again from the explicit separation of the field into condensed and noncondensed parts, precluding its use in strongly nonequilibrium and turbulent regimes of Bose-field dynamics.

2.3 Superfluidity and vortices

2.3.1 Superfluidity

The term ‘superfluid’ is used to describe a complex of phenomena that can arise in *quantum* fluids, that distinguish them from ordinary (or *normal*) fluids [162]. The properties of a superfluid are fundamentally related to the presence of a single macroscopic quantum wave function in the fluid, and the nature of the excitations of the fluid relative to this quantum degenerate ‘ground state’. Some of the characteristic features of superfluidity are a direct result of the flow being described by a quantum wave function, and are foreshadowed by results in single-particle quantum mechanics; for example, the motion of electrons in a rotating molecule ‘slips’ with respect to the rigid rotation of the nuclei [262]. Superfluid many-body systems (such as superfluid helium) exhibit similar behaviour [128]. Superfluidity is characterised by the appearance of a *superfluid velocity* \mathbf{v}_s , for which, from the transformation properties of the order parameter associated with condensation in the system, we deduce [207]

$$\mathbf{v}_s = \frac{\hbar}{m} \nabla S, \quad (2.54)$$

where $S(\mathbf{x}, t)$ is the *quantum phase* of the order parameter (i.e., of the condensate). It is important to note however that while the phase of the condensate determines the velocity of the superflow, the mass current of the superflow is not, in general, directly related to the density of the condensate [207]; i.e., Bose-Einstein condensation and superfluidity

⁴We note that in the standard symmetry-breaking approach one is free to *choose* the excitations (and thus $\hat{\delta}(\mathbf{x})$) to be orthogonal to the condensate [184], but this is not required by the derivation.

are intimately related but distinct concepts. A conspicuous example of this is provided by superfluid ^4He , in which the condensate fraction is a mere 8% at zero temperature, while the superfluid fraction is unity (see, e.g., [132]). However in the limit of complete Bose condensation in weakly interacting gases the superfluid and condensate densities are approximately equal [163].

Stability of superflow

A crucial aspect of bulk superfluids is the stability of their flow, arising from the difficulty in creating excitations in the fluid, which is a direct consequence of the spectrum of excitations about the ground state. This can be understood already from the Bogoliubov spectrum of a homogeneous Bose gas. In such a scenario the condensate forms in the $k = 0$ (i.e., constant) plane-wave mode, and the excitations are mixtures of plane waves, with the spectrum [47, 207]

$$\begin{aligned}\epsilon_k &= \left[\frac{\hbar^2 k^2}{2m} \left(\frac{\hbar^2 k^2}{2m} + 2U_0 n \right) \right]^{\frac{1}{2}} \\ &\approx \hbar k \sqrt{\frac{U_0 n}{m}} \quad \text{for} \quad \frac{\hbar^2 k^2}{2m} \ll 2U_0 n.\end{aligned}\tag{2.55}$$

At long wavelengths, the spectrum, which is linear in k is that of *sound waves*, i.e., the wave velocity $v_k = d\omega_k/dk = \sqrt{U_0 n/m} \equiv c_s$ is independent of the wave number k . Simple energy and momentum conservation arguments [166, 47, 203] then show that an impurity moving through the fluid can not excite quasiparticle excitations of the condensate if its velocity $v < c_s$, which is known as *Landau's criterion* for superflow stability. The motion of such an object through the fluid is thus not resisted by the superfluid component. By a Galilean transformation one obtains the archetypal property of superfluidity: dissipationless flow of the fluid through a capillary [207]. In general finite-temperature scenarios there is of course also a *normal* component of the fluid, which can exert a drag force on the impurity or capillary wall. In bulk scenarios, the two fluids also exert no friction on each other [207], and can be described by a *two-fluid model* [158]. This however breaks down in inhomogeneous superfluids such as the trapped condensates we consider in this thesis, where the superfluid can be excited at its ‘surface’ (where the sound velocity vanishes with the superfluid density), and even in the homogeneous case in the presence of defects in the superfluid, such as quantum *vortices*, which we discuss in section 2.3.2.

Landau criterion for surface modes of a trapped condensate

A useful result for our investigations of rotating systems in this thesis is obtained by transposing the Landau criterion for the stability of superflow to the surface modes (section 2.2.2) of a Bose condensate. Making this transposition we obtain the critical angular frequency [59]

$$\Omega_L = \min \left(\frac{\omega_m}{m} \right), \quad (2.56)$$

where ω_m is the frequency of a surface mode with angular-momentum projection m . It is important to note the physical meaning of this result: Ω_L is the minimum rotation frequency at which a surface mode is shifted (by the inertial term $-\Omega L_z$ in the rotating-frame Hamiltonian) to a negative value. That is, in a frame rotating at an angular frequency $\Omega > \Omega_L$, the surface of the condensate is thermodynamically unstable with respect to the growth of surface modes. As we will discuss in chapter 8, this result is confirmed by calculations which explicitly include the dissipative effects of the thermal cloud [266, 200].

2.3.2 Vortices

Given a condensate wave function $\phi(\mathbf{x}) = |\phi(\mathbf{x})|e^{iS(\mathbf{x},t)}$, the velocity of the superflow is given by equation (2.54). An immediate consequence of this is that the superflow is irrotational, i.e., we must have $\nabla \times \mathbf{v}_s = \mathbf{0}$. Stokes' theorem then implies that the integral of the velocity $\oint_C \mathbf{v}_s \cdot d\boldsymbol{\ell} = 0$ around any closed path C which encloses a simply-connected domain. If the enclosed domain is not simply connected, then this result does not hold, and the only relevant constraint is that imposed by the single-valued nature of the phase, which is however only defined up to the periodicity (2π) of the complex phase. We therefore find the Feynman-Onsager quantisation condition

$$\oint_C \mathbf{v}_s \cdot d\boldsymbol{\ell} = n\kappa, \quad (2.57)$$

where $\kappa \equiv h/m$ is the *quantum of circulation*, and n is some integer. A condensate which is otherwise simply connected can therefore support a circulation by admitting a node about which the phase of the wave function wraps continuously from 0 to $2n\pi$. In two dimensions such a node is called a point vortex, in three dimensions the node has finite extent and forms a *vortex line* (or vortex filament). In this thesis, we will consider vortices only in quasi-two-dimensional systems, and we will therefore deal only with point vortices. The fluid flow remains irrotational everywhere, as the phase and thus velocity




Figure 2.1: (a) Motion of a vortex with respect to the superfluid background: the vortex velocity is perpendicular to the imposed external force (see equation (2.58)). (b) Motion of two like vortices. Each vortex is advected by the flow of the other, and they therefore orbit about their geometric centre. (c) Effect of the longitudinal friction force on a vortex: in the frame of the normal fluid, the vortex drifts at an angle to the background superflow.

are meaningless at the node, where the wave function vanishes, however, the vortex can be thought of as a point ‘charge’ of rotation.

Dynamics of vortices

In practice, vortices with circulation $|n| > 1$ are energetically unstable to decay into n singly charged vortices [75]. Similarly, considering the energy of a collection of vortex lines, one can show that this energy takes the same form as the electrostatic energy of a collection of charged filaments [243], i.e., vortices of opposite sign exert an attractive force on one another, and vortices of like sign exert a repelling force on each other. However, by Galilean invariance, one finds [86] that vortices are carried along by the local flow of the superfluid, and thus by the flow of other vortices. This somewhat counterintuitive relation between the force between the vortices and their motion is an example of the peculiar dynamics obeyed by vortices. Hydrodynamic considerations yield the result that vortex moving (with velocity \mathbf{v}_1) against a background superflow (velocity \mathbf{v}_s) experiences the *Magnus force* [75] $\mathbf{F}_m = \rho_s \boldsymbol{\kappa} \times (\mathbf{v}_1 - \mathbf{v}_s)$, where ρ_s is the density of the superfluid, and $\boldsymbol{\kappa}$ is a vector with magnitude κ directed along the axis of circulation of the vortex. As a result, the motion of the vortex subject to an applied external force \mathbf{F}_{ext} is governed by the law [244]

$$\rho_s [(\mathbf{v}_1 - \mathbf{v}_s) \times \boldsymbol{\kappa}] = \mathbf{F}_{\text{ext}}. \quad (2.58)$$

Equation (2.58) shows that, in contrast to Newtonian dynamics in which the acceleration of an object is proportional to the force applied to it, the *velocity* of a vortex relative to

the superflow is proportional to the force acting on it [243]⁵, as illustrated in figure 2.1(a). Consequently two like vortices which energetically ‘repel’ each other tend to orbit about their geometric centre (in the absence of dissipation), as illustrated in figure 2.1(b). As a result of this energetic repulsion and tendency to orbit about one another, a superfluid to which a large rotation is imparted acquires a large array of singly charged vortices, and the lowest-energy configuration of such an array is a crystalline vortex *lattice* (see section 2.3.3).

It is clear from this discussion that there is an analogy between vortices in a superfluid and charged particles (or filaments) moving in a background magnetic field. Indeed there is a strong analogy between vortices in a (uniform) two-dimensional superfluid and a (2+1)-dimensional electrodynamics [15], with phonons playing the role of electromagnetic radiation. Electrostatic and electrodynamic analogies are therefore a very powerful tool in understanding the behaviour of vortices in superfluid systems.

Vortices at finite temperature

At a finite temperature, we must include the effects of thermal interactions on vortex motion. The fundamental process of interest is the scattering of excitations by the vortex core [244], which produces a conventional (longitudinal) drag force on the vortex $F_{\parallel} = -D(\mathbf{v}_1 - \mathbf{v}_n)$, where \mathbf{v}_n is the velocity of the normal component, and D is a (temperature dependent) coefficient characterising the strength of the friction. As one might expect, this frictional force opposes the motion of the vortex, but the corresponding motion is of course determined by the Magnus law (2.58). For example, a vortex moving against a stationary normal fluid $\mathbf{v}_n = \mathbf{0}$ moves at some angle to the background superflow, as illustrated in figure 2.1(c). The existence and magnitude of a second *transverse* force component due to the scattering of excitations remains a controversial issue [244, 13, 94]. Such a force component would be nondissipative, and numerical calculations [22, 140] suggest that if this force is indeed nonzero, it is certainly small.

The interaction of vortices with the normal flow further complicates the dynamics of a realistic superfluid system, and the additional effects are included at a phenomenological level in the extended Hall-Vinen-Bekharevich-Khalatnikov theory [125]. This model applies in the limit that the spacing between vortex lines is small, and the vorticity of the superfluid (and its interaction with the normal fluid) is treated as a continuum property of the superfluid.

⁵This non-Newtonian dynamical law has some amusing consequences. For example, a system of three vortices in an unbounded flow in 2D is integrable, the motion only becoming chaotic upon the addition of a *fourth* vortex [14].

2.3.3 Vortices in Bose condensates

Central-vortex states

A condensate in a cylindrically symmetric harmonic trap with a single, *central*, singly charged vortex parallel to the z axis has a Gross-Pitaevskii wave function of form

$$\psi(\mathbf{x}) = e^{i\theta} |\psi(r_\perp, z)|, \quad (2.59)$$

where $r_\perp = \sqrt{x^2 + y^2}$ is the radial distance from the z axis, and θ is the azimuthal angle around this axis. In the noninteracting limit, this wave function is of form

$$\psi_1(\mathbf{x}) = \frac{N}{\pi^{3/4} \sqrt{d_\perp^2 d_z}} e^{i\theta} r e^{-\frac{1}{2}(r_\perp^2/d_\perp^2 + z^2/d_z^2)}, \quad (2.60)$$

where $d_\perp = \sqrt{\hbar/m\omega_\perp}$ and $d_z = \sqrt{\hbar/m\omega_z}$ are the oscillator lengths corresponding to the confinement in the r_\perp and z directions, respectively. In an inertial (laboratory) frame, this wave function has higher energy than that of the condensate consisting of N atoms in the ground state of the trap, $\psi_0(\mathbf{x}) = [N/(\pi^{3/2} d_\perp^2 d_z)] \exp\{-\frac{1}{2}(r_\perp^2/d_\perp^2 + z^2/d_z^2)\}$; i.e., it is an excited state. However, in a rotating frame the energy of the vortex will be lowered due to its circulation. We thus consider the relative energies of the two states in a frame rotating with angular velocity Ω about the z axis. In such a frame the GP energy is given by

$$E_\Omega[\psi] = E - \Omega \langle L_z \rangle = \int d\mathbf{x} \psi^* \left[-\frac{\hbar^2}{2m} \nabla^2 + V_{\text{ext}} - \Omega L_z + \frac{U_0}{2} |\psi|^2 \right] \psi. \quad (2.61)$$

As the vortex-free ground state $\psi_0(\mathbf{x})$ carries no circulation, its energy E_0 is the same in all rotating frames, while the energy E_1 of the vortex state $\psi_1(\mathbf{x})$ is reduced from its lab-frame value by $N\hbar\Omega$ in the frame rotating at angular velocity Ω . As such there is a *critical* angular velocity at which the energy of the vortex state becomes equal to that of the vortex-free state, and for higher angular velocities, the vortex state is a lower energy state than the vortex-free state (see figure 2.2(a)). However, in the noninteracting limit one can easily see that, as the lab-frame energy of the vortex state is $E_1 = E_0 + N\hbar\omega_\perp$, this crossover does not occur in the noninteracting case until $\Omega = \omega_\perp$, in which limit the trapping potential is completely cancelled by the centrifugal force resulting from the rotation [87].

In the opposite Thomas-Fermi limit of *dominant* interactions, the vortical condensate has density approximately of the form [87]

$$n(r_\perp, z) \approx \frac{\mu}{U_0} \left(1 - \frac{\eta^2}{r_\perp^2} - \frac{r_\perp^2}{R_\perp^2} - \frac{z^2}{R_z^2} \right) \Theta \left(1 - \frac{\eta^2}{r_\perp^2} - \frac{r_\perp^2}{R_\perp^2} - \frac{z^2}{R_z^2} \right), \quad (2.62)$$




Figure 2.2: (a) Rotating-frame energies of a generic vortex-free ground state and the corresponding central-vortex state. The energy of the vortex state decreases with increasing frame angular velocity, and at a critical frequency Ω_c its energy drops below that of the vortex-free state. (b) Energy of a vortex state as a function of the radial displacement of the vortex, evaluated in various rotating frames. Cases are: (i) $\Omega = 0$, thermodynamically unstable state; (ii) $\Omega = \Omega_m$, onset of vortex metastability; (iii) $\Omega = \Omega_c$ onset of thermodynamic stability of the vortex; (iv) $\Omega = (3/2)\Omega_c$, stable state with energy barrier.

where R_\perp, R_z are the appropriate Thomas-Fermi radii, and η is the healing length. As the state is a singly charged on-axis vortex, it carries angular momentum $L_v = N\hbar$, and it can be shown that the lab-frame energy of this state is [172]

$$E_v = \frac{4\pi}{3} \frac{n_0 \hbar^2 R_z}{m} \ln \left(\frac{0.67 R_\perp}{\eta} \right), \quad (2.63)$$

where n_0 is the peak density of the corresponding vortex-free TF state, with healing length η . Comparing the vortex-state energy equation (2.63) with that of the corresponding vortex-free state, we find that the critical frequency at which the vortex state becomes the ground state is [172]

$$\Omega_c = \frac{5}{2} \frac{\hbar^2}{m R_\perp^2} \ln \left(\frac{0.67 R_\perp}{\eta} \right). \quad (2.64)$$

We note that the vortex-state energy equation (2.63) is obtained by assuming the healing length η as short-range cutoff in the calculation of the kinetic energy due to the vortex flow pattern (cf. [111]). Numerical solution of the full Gross-Pitaevskii equation reveals that the true energy contribution of the vortex is somewhat higher [87]. Nevertheless equation (2.64) agrees well with numerically exact results [61] for the critical frequency in the limit of large condensates [172].

Off-axis vortices

So far we have shown the energy of the vortex at the centre of the trap is lower than that of the vortex-free state in a frame rotating above some critical angular velocity Ω_c . Here we consider the approximate energy of the vortex state as a function of the displacement of the vortex from the trap centre, thus continuously connecting the central-vortex and vortex-free limits. A Thomas-Fermi limit calculation for the energy of state with vortex displaced off-axis by r_v , in a frame rotating at angular velocity Ω , leads to the result [87, 153]

$$E_v(r_v; \Omega) = \left[1 - \left(\frac{r_v}{R_\perp} \right)^2 \right]^{\frac{3}{2}} \left\{ E_v(0; 0) - N\hbar\Omega \left[1 - \left(\frac{r_v}{R_\perp} \right)^2 \right] \right\}, \quad (2.65)$$

where $E_v(0; 0)$ is the lab-frame energy of the central-vortex state given by equation (2.63). In figure 2.2(b) we plot the results of this energy as a function of the vortex displacement, for several values of the rotating-frame angular velocity Ω . In the laboratory frame (case (i)), the energy of the state decreases monotonically with increasing vortex displacement, showing that the central vortex is truly thermodynamically unstable in this regime; in the presence of weak dissipation in this frame (i.e., in the presence of stationary thermal cloud), the dissipation serves to expel the vortex from the condensate. In case (ii) we illustrate the form of the energy in a frame rotating at frequency

$$\Omega_m = \frac{3}{2} \frac{\hbar}{mR_\perp^2} \ln \left(\frac{0.67R_\perp}{\eta} \right), \quad (2.66)$$

at which the curvature of $E_v(r_v)$ at $r_v = 0$ changes sign. For greater angular velocities $\Omega > \Omega_m$ the function $E_v(r_v)$ is convex at $r_v = 0$, and thus in the presence of weak dissipation in this frame the vortex is pushed back towards the centre of the trap. As the vortex energy at the trap centre is still higher than the energy of the vortex-free state, this frequency marks the onset of *metastability* of the central-vortex state.

In case (iii) we indicate the curve corresponding to the frequency Ω_c given by equation (2.64) (note that $\Omega_c = (5/3)\Omega_m$). As we have discussed, beyond this frequency the energy of the central-vortex state is lower than that of the vortex-free state, i.e., the vortex state is the global energy minimum. However, even for significantly higher frequencies (case (iv), $\Omega = (3/2)\Omega_c$), we note the presence of an *energy barrier* that must be surmounted by the vortex in order for it to migrate from the outside of the condensate to the centre. Thus, even in the presence of weak dissipation, above the frequency Ω_c additional mechanisms are necessary in order for the vortex to migrate to the trap centre, as we discuss in chapters 8 and 9.

The anomalous mode

An *anomalous mode* is a Bogoliubov excitation of the GP state with *negative energy*, the appearance of which signals the thermodynamical instability of the condensate (section 2.2.3). As a central single-vortex state is thermodynamically unstable in the lab frame (figure 2.2(b)), such an anomalous mode arises in the Bogoliubov expansion about such a state [73]. In the noninteracting limit, this mode simply reduces to the ground state of the trap [223]. For the interacting case, in the Bogoliubov description, the u and v components of this anomalous mode have circulation 0 and 4π respectively, indicating that the mode has angular momentum [181] $q = -1$ with respect to the central-vortex state. The energy of this state (relative to that of the condensate) therefore increases with increasing angular velocity of the rotating frame, and reaches zero at the metastability frequency Ω_m , as we expect. The energy of the anomalous mode of the central-vortex state therefore determines the frequency of precession of a vortex in the limit of *infinitesimal* displacement of the vortex from the trap axis [87]. More generally, an anomalous mode appears for a *finitely* off-axis vortex, and has an excitation energy related to the precession frequency of the vortex [168]. However, this relationship has been shown not to hold in the extension to finite-temperature HFB theories [138, 264].

Vortex lattices

As multiply charged vortices are unstable, the equilibrium configuration of a macroscopically rotating superfluid is expected to contain a proliferation of singly charged vortices. In the limit of large angular momentum, we expect such a rotating superfluid to exhibit a *coarse-grained* velocity profile which mimics that of solid-body rotation. The latter is given by a velocity $\mathbf{v}_{sb} = \boldsymbol{\Omega} \times \mathbf{x}$ at position \mathbf{x} , and yields a constant vorticity $\nabla \times \mathbf{v}_{sb} = 2\boldsymbol{\Omega}$. This prompted Feynman [90] to estimate the density of vortices in a macroscopically rotating fluid as the ratio of this vorticity to the elemental vorticity κ , yielding

$$n_v = \frac{2\Omega}{\kappa} = \frac{m\Omega}{\pi\hbar}, \quad (2.67)$$

which is referred to as *Feynman's rule* [75]. Calculations performed by Tkachenko show that the minimum energy configuration is one in which the vortex lattice is *triangular* [253], and in the literature such lattices are known as Abrikosov lattices, named for Abrikosov who predicted the analogous structure for magnetic flux lines in type-II superconductors [4]. We note, however, that calculations performed by Campbell and Ziff [43] indicated that the rotational velocity field in fact leads to significant deviations from this ideal triangular structure in uniform, cylindrically confined superfluids. It is at first some-

what surprising then that the lattices formed in dilute Bose condensates (e.g., [2, 124]) are so strikingly regular, especially considering that one might expect the vortex density to be lower at the centre of the trap, where the superfluid density and thus the kinetic energy associated with vortices is highest. However, Sheehy and Radzihovsky [233, 234] have shown in a careful analysis that this energetic effect is almost exactly balanced by the spatial dependence of the vortex ‘chemical potential’ ($\propto \Omega \rho_s$). They obtain a general result for the vortex density, which in the TF limit (for isotropic or oblate trapping) reduces to the radius-dependent result

$$n_v(r) \approx \frac{m\Omega}{\pi\hbar} - \frac{1}{2\pi} \frac{R^2}{(R^2 - r^2)^2} \ln \left(\frac{e^{-1}\hbar}{\eta^2 \Omega m} \right), \quad (2.68)$$

where R is the Thomas-Fermi radius, and e is the base of the natural logarithm.

Chapter 3

Classical-field theory

In this chapter, we introduce the classical-field formalism we use throughout this thesis. The term ‘classical-field theory’ in fact refers to a *family* of theoretical and computational formalisms for the degenerate Bose gas, with generally different regimes of applicability and utility, which have been recently reviewed in [28]¹. However, they all share the feature of allowing for the description of many-body physics beyond the reach of a simple Gross-Pitaevskii treatment, while retaining the computational tractability of a description in terms of a single-particle state space. In this thesis, we focus on the Hamiltonian (i.e., deterministic) dynamics of classical fields described by the *projected Gross-Pitaevskii equation*, and in particular on the description of *thermal* Bose fields using this formalism, in both equilibrium and nonequilibrium scenarios. This chapter serves to explain the physical content and formulation of this particular methodology, and to place it in the broader context of classical-field methods for the degenerate Bose gas.

We begin this chapter in section 3.1 with a heuristic introduction to the basic notions of the classical-field treatment, and explain its validity and utility for the finite-temperature scenarios we will consider in this thesis. In section 3.2 we motivate the low-energy *projection* of the field that is a central aspect of the method, and present the derivation of the projected Gross-Pitaevskii equation from the second-quantised Hamiltonian for the dilute Bose gas previously introduced in chapter 2. In section 3.3 we describe the ergodic character of solutions of the projected GP equation, and discuss the *thermodynamics* of the classical field that emerge dynamically from the field trajectories. Finally in section 3.4 we introduce the closely related *truncated-Wigner* approximation for the degenerate Bose gas, and discuss its relation and relevance to the classical-field formalism we employ in this thesis.

¹We note that the authors of [28] coined the term ‘c-field’ to describe the collection of classical-field methods discussed in that work, so as to avoid the potential confusion associated with the phrase ‘classical field’, which can give the misleading impression that these methods are not quantum mechanical.

3.1 Nature of the classical-field approximation

At its simplest level, classical-field theory results from the replacement of a second-quantised (Bose) field Hamiltonian with an analogous classical field Hamiltonian. This is motivated by the observation that in regimes of thermal behavior, the dominant characteristic of the atomic Bose-field system is its multimode, self-interacting nature, rather than the quantised nature of the atomic field. Indeed, under conditions of high mode occupation, field commutators become relatively unimportant, and a satisfactory description of the bosonic field can be obtained from a classical-field model. It is important to note that while this replacement is superficially the same as the replacement of the field operator $\hat{\Psi}(\mathbf{x})$ by a classical (c-number) field in the Gross-Pitaevskii approximation (section 2.2.2) the interpretation is very different. For example, the GP wave function $\phi(\mathbf{x})$ describes a condensate, which carries no entropy, while by contrast, the classical field in our models is typically a strongly fluctuating field, and through its fluctuations, describes *thermal* behaviour of the field which is *by definition* outside the regime of validity of the GPE.

In general, we will think of the classical-field descriptions of the systems we study as approximations to the full second-quantised field theory. However, we can also view the classical-field system as the ‘original’ classical-dynamical system, from which the quantum field theory is obtained by quantising *once* [202]. Taking this viewpoint, it makes sense to characterise the behaviour of this classical-field system first, before layering on the complexity of field quantisation.

This viewpoint also lends us more insight into the nature of the classical-field approximation, when we consider an alternative formulation of the quantum field theory: the *path integral* formulation [89, 275]. In this formulation, the transition probability for the field to evolve from a configuration $\psi_1(\mathbf{x}, t_1)$ at time t_1 to another configuration $\psi_2(\mathbf{x}, t_2)$ at a later time t_2 is given by a sum (path integral) over phase factors $e^{iS[\psi(\mathbf{x}, t)]/\hbar}$, where $S[\psi(\mathbf{x}, t)]$ is the *action* evaluated for each conceivable path between the end configurations (figures 3.1(a-b)). In this way, quantum fluctuations (the physical content of the commutation relations (2.4)) are built into the theory. The classical field trajectory is regained in the limit that the contribution of these quantum fluctuations to the path integral is small², in which case the path of least action (i.e., the classical trajectory) dominates the integral (figure 3.1(a)). In the situations we consider in this thesis, we evolve a *single* classical trajectory of the field, which exhibits strong fluctuations (figure 3.1(c)) driven

²In single-particle quantum mechanics this limit is of course $\hbar \rightarrow 0$. However, our classical field equation contains an explicit \hbar , so the limit here is more accurately regarded as that of infinite particle number [211].

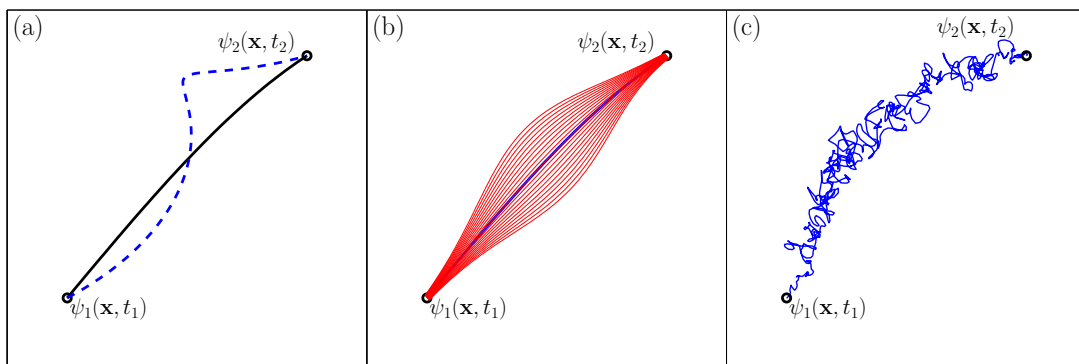


Figure 3.1: (a) Evolution of a simple classical field in the Lagrangian viewpoint: the physical trajectory (solid line) is that for which the action is minimised. (b) Path integral interpretation of quantum field theory: all possible paths contribute to the probability of the transition from ψ_1 to ψ_2 ; strongly nonclassical trajectories interfere destructively. (c) Self-interacting classical field: the path of least action is itself strongly fluctuating, and samples many different field configurations during the classical evolution.

by the kinetics of its own self-interaction. This strongly fluctuating character is more significant than the ‘spread’ of contributing paths included in the quantum theory, and so we do not consider the contributions of these other paths. We thus obtain a *nonperturbative* description of the field dynamics, at the cost of sacrificing the Bose commutation relations (2.4) and the results which follow from them (such as the exact Bose-Einstein thermodynamic distribution in equilibrium). Another important feature of the fluctuating classical trajectory of the field illustrated in figure 3.1(c) is that we can extract meaningful statistics by averaging over the fluctuations *along* the path of the field (rather than over multiple neighbouring paths). In equilibrium this reflects the fact that the classical field is *ergodic*, and samples from the appropriate thermal distribution over time, as we discuss in section 3.3. However, this also provides a representation of thermal fluctuations of the field in scenarios away from equilibrium. A strength of the classical-field method is that it implicitly includes the effect of these thermal fluctuations on the *dynamics* of the field, thus providing a formalism for modelling *finite-temperature nonequilibrium* behaviour of the atomic Bose field.

3.2 Projection

3.2.1 Motivation

A crucial aspect of the classical-field formulation we use in this work is the *projection* of the classical field onto a low-energy region. In practice we simulate a field which differs

from the second-quantised Bose field $\hat{\Psi}(\mathbf{x})$ not only in that its quantum-operator character is neglected, but also in that it comprises an expansion over a *finite* number of single-particle basis modes, which are typically chosen as the lowest-energy eigenmodes of some single-particle Hamiltonian. In this section we briefly outline some of the motivations for the restriction of the field in this manner, before introducing the projection formally in section 3.2.2.

Numerical implementation of field theory

The first motivation is a practical one: a true ‘field theory’ is a theory of the dynamics of a system with an infinite number of degrees of freedom, i.e., the field is obtained as the continuum limit of a (classical or quantum) dynamical system with a discrete set of dynamical variables [252]. Obviously, it is not in general possible to describe such a system with finite computational resources. In practice, every computational realisation of field theory must be defined in terms of only some subset of the true degrees of freedom of the field theory, and so we are compelled to consider carefully which degrees of freedom we include in our numerical model and which we discard. This is doubly important as naive discretisations of the theory can introduce serious pathologies into the dynamics, as we discuss further in section 4.3, when we compare the projected formalism to other common discretised formulations of the Gross-Pitaevskii theory.

Long-wavelength nature of the theory

The second motivation is that the quantum field theory corresponding to the Hamiltonian (2.20) can only be regarded as an *effective field theory*, describing the behaviour of the atomic field on wavelengths which are long compared to some *cutoff* length scale. This is evidenced by the use of a contact (*s*-wave) interaction potential, which implicitly assumes we are only describing the dynamics of wavelengths of the field large compared to the range r_0 of the interatomic potential [104], such that we can neglect the true structure of this potential. Moreover, the contact potential is independent of the relative momentum, and thus describes the same interaction strength for all collisions, whereas physically the efficiency of scattering by the true inter-atomic potential is curtailed dramatically for relative momenta $k > 1/a_s$. Note that this issue is already present in the second-quantised Hamiltonian (2.6) from which we derive the Gross-Pitaevskii equation, but does not lead to problems in that case as we consider only the condensate mode, which is essentially the ‘longest-wavelength’ mode³, and for which the *s*-wave assumption holds true. It does, however, lead to ultraviolet divergences in the self-consistent mean-field theories which

³In particular, (repulsive) interactions act to ‘smooth’ the GP mode so that it is slowly-varying in space, and contains only long-wavelength components.

generalise the GP description to finite temperatures. Morgan [183] has shown that by formally restricting the description of the atomic field to its long-wavelength components, the correct description of the interparticle interaction experienced by this field appears as precisely the ‘contact’ potential used in the continuum theory, but this interaction potential is now formally regular due to the long-wavelength nature of the field, and this is the scenario we consider in this thesis. Formally the scattering coefficient is then a cutoff-dependent parameter [183, 239, 28], however in practice the correction due to the cutoff is small [183, 190], and we can simply use the experimentally measured value for a_s in our models.

Validity of the classical-field approximation

The third motivation is that by choosing the cutoff at an appropriate height in energy, it is possible (in some scenarios) to ensure that all single-particle modes in the classical-field description are highly occupied, i.e., that their classical occupations are $\gtrsim 1$, dominating the neglected ‘vacuum occupation’ of the modes. In such a regime the vacuum fluctuations have little effect on the dynamics and the classical-field approximation is formally justified. At thermal equilibrium of the classical field in this limit, the thermodynamics of the classical field can be identified with the high-occupation limit of the true Bose-field system, and highly accurate [65] estimates for the correlations and thermodynamics of the quantum many-body system can be obtained by including a description of *above-cutoff* atoms into the theory. However, in general scenarios (away from equilibrium, for example), this restriction of the domain of the classical field so that all modes are highly occupied may not be practical, or even possible. In this thesis, we focus on the *intrinsic* phenomenology of the classical field system, and so we will frequently relax this requirement. As a result, our results are not expected to be quantitatively accurate predictions for the true bosonic field theory. Nevertheless, we find that the behaviour of the classical field yields great insight into the behaviour of the atomic Bose field, in scenarios which are essentially inaccessible to other contemporary finite-temperature formalisms, built on alternative approximations to the underlying many-body theory.

Expression in single-particle energy levels

As we have already noted, the way in which the atomic interaction is represented in our model requires us to restrict the field so as to exclude high-momentum components. However, in the systems we consider in this thesis, the atomic field is *confined* by an external trapping potential, which breaks the translational symmetry of the system. Momentum is

thus no longer a conserved quantity of the system, and so it is not appropriate to impose the classical-field cutoff in such terms.

In typical conditions, we expect all single-particle modes of the classical field below the cutoff to be populated, including those which ‘border’ the cutoff. Due to the complexity of the interacting field theory, requisite conditions for the cutoff are best formulated by considering thermal equilibrium (or quasi-equilibrium) scenarios. In such cases, we expect the transfer of population to obey the principle of detailed balance [221]: the transfer of population from a mode j to a mode k is exactly balanced by the transfer of population from k to j . By carefully applying a cutoff we enforce this behaviour for transfer between modes on opposite sides of the cutoff: detailed balance is ensured in the trivial sense that the transfer in both directions is identically zero. Applied correctly, this cutoff should then have no effect on the dynamics, at least close to equilibrium. However, it is important that the cutoff is defined in a way which respects the symmetries of the system as otherwise the resulting equilibrium of the below-cutoff field will not be, in general, consistent with an equilibrium of the full system. In such cases the effective detailed-balance condition resulting from the cutoff will be spurious. For this reason, the cutoff should be defined in terms of the *energy* of the system, so that the cut can be made with minimal disruption to the dynamics and thermodynamics of the below-cutoff field. Equivalently, we require that the constant-energy surface of the cutoff system most closely matches a constant-energy surface of the full interacting system.

Of course, in general the surface of constant energy for the interacting system will be a very complicated function of the field variables, which is impossible to calculate. However, we do know that for an equilibrium system containing a condensate, the energy of the system can be approximately diagonalised by the Bogoliubov transformation (section 2.2.3). Moreover, we know that at high energies the Bogoliubov modes return to single particle states, i.e. $v_k \rightarrow 0$ as $\epsilon_k \rightarrow \infty$ [60], and in this limit the u_k approach the energy eigenstates of the appropriate single-particle Hamiltonian (with shifted energies). In fact, in harmonic trapping, for energies $\gtrsim 3\mu$, where μ is the chemical potential of the condensate, the energies of the Bogoliubov modes and those of the single-particle energy eigenstates agree to within about 3% [105]. At such energies the presence of the condensate simply results in a small upwards shift of the single-particle mode energies. Therefore, effecting the cutoff in terms of the single-particle energy in the trapping potential, we can enforce a cutoff which is *approximately* diagonal in the energy of the interacting system. In practice, this cutoff is implemented by an explicit *projection* onto single-particle eigenstates with energies below the cutoff.




Figure 3.2: Schematic view of the coherent region (condensate band) **C** and incoherent region (noncondensate band) **I** for a harmonic trap.

3.2.2 Projected classical-field theory

The classical-field energy functional

In our projected classical-field theory, we introduce a cutoff energy E_R to separate the system into a low-energy *coherent region* (**C**) consisting of single-particle modes energies satisfying $\epsilon_n < E_R$, and a high-energy *incoherent region* (**I**) containing the remaining modes, as indicated in figure 3.2. These regions are also sometimes referred to as the condensate and noncondensate *bands*, respectively [102, 104, 37]. Where possible, it is most convenient to carry out the separation in the basis which diagonalises the single-particle Hamiltonian of the system, since at sufficiently high energies ($\epsilon_n \sim E_R$ [105, 104]) the many-body Hamiltonian (2.20) is approximately diagonal in this basis, as discussed in section 3.2.1. For certain potentials, an appropriate numerical quadrature method exists which allows for the efficient integration of the resulting equation of motion. More generally, the single-particle Hamiltonian is divided into a *base* Hamiltonian and a remaining potential part which is possibly time dependent

$$H_{\text{sp}} = H_0 + \delta V(\mathbf{x}, t). \quad (3.1)$$

The total potential in general comprises a contribution $V_0(\mathbf{x})$ from the base Hamiltonian and the perturbation potential, i.e. $V(\mathbf{x}, t) = V_0(\mathbf{x}) + \delta V(\mathbf{x}, t)$, while the base Hamiltonian may contain inertial terms associated with the transformation to a rotating frame.

The validity of implementing the cutoff in terms of the single-particle energy corre-

Chapter 3. Classical-field theory

sponding to H_0 requires that $\delta V(\mathbf{x}, t)$ be a small perturbation to H_0 (so that, e.g. the discrepancy between the constant-energy surfaces of H_0 and H_{sp} at the cutoff is of the same order as that between the single-particle and Bogoliubov Hamiltonians). We then introduce the projected field operator

$$\hat{\psi}(\mathbf{x}) \equiv \sum_{n \in \mathbf{C}} \hat{a}_n \phi_n(\mathbf{x}) = \mathcal{P} \hat{\Psi}(\mathbf{x}), \quad (3.2)$$

where the projector is defined by

$$\mathcal{P} f(\mathbf{x}) \equiv \sum_{n \in \mathbf{C}} \phi_n(\mathbf{x}) \int d\mathbf{x}' \phi_n^*(\mathbf{x}') f(\mathbf{x}'). \quad (3.3)$$

The summation is over all modes satisfying $\epsilon_n \leq E_R$, where $H_0 \phi_n(\mathbf{x}) = \epsilon_n \phi_n(\mathbf{x})$ defines the modes in terms of the base Hamiltonian. The number of single-particle modes included in the coherent region is the *condensate-band multiplicity*, denoted by \mathcal{M} . The projected field operator thus obeys the commutation relation

$$[\hat{\psi}(\mathbf{x}), \hat{\psi}^\dagger(\mathbf{x}')] = \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x}'), \quad (3.4)$$

where the *projected delta function*

$$\delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x}') = \sum_{n \in \mathbf{C}} \phi_n(\mathbf{x}) \phi_n^*(\mathbf{x}'). \quad (3.5)$$

In the s -wave limit, the effective low energy Hamiltonian for the system is then given by

$$H_{\text{eff}} = \int d\mathbf{x} \hat{\psi}^\dagger(\mathbf{x}) H_{\text{sp}} \hat{\psi}(\mathbf{x}) + \frac{U_0}{2} \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}). \quad (3.6)$$

We now obtain a classical theory by making the replacement

$$\hat{\psi}(\mathbf{x}) \rightarrow \psi(\mathbf{x}) \equiv \sum_{n \in \mathbf{C}} \alpha_n \phi_n(\mathbf{x}), \quad (3.7)$$

in H_{eff} which yields the classical-field *energy functional*

$$H_{\text{CF}} = \int d\mathbf{x} \psi^*(\mathbf{x}) H_{\text{sp}} \psi(\mathbf{x}) + \frac{U_0}{2} |\psi(\mathbf{x})|^4 \equiv E[\psi]. \quad (3.8)$$

Classical mechanics interpretation

The energy functional (3.8) is in fact a classical Hamiltonian, from which we can derive the evolution of the projected field $\psi(\mathbf{x}, t)$. To see this, we write H_{CF} explicitly in terms

of the classical-field amplitudes $\{\alpha_j, \alpha_j^*\}$, and obtain

$$\begin{aligned} H_{\text{CF}} = & \sum_{n \in \mathbf{C}} \epsilon_n \alpha_n^* \alpha_n + \sum_{mn \in \mathbf{C}} \int d\mathbf{x} \phi_m^*(\mathbf{x}) \delta V(\mathbf{x}, t) \phi_m(\mathbf{x}) \alpha_m^* \alpha_n \\ & + U_0 \sum_{mnpq \in \mathbf{C}} \int d\mathbf{x} \phi_m^*(\mathbf{x}) \phi_n^*(\mathbf{x}) \phi_p(\mathbf{x}) \phi_q(\mathbf{x}) \alpha_m^* \alpha_n^* \alpha_p \alpha_q. \end{aligned} \quad (3.9)$$

By analogy to the relation between the quantum ladder operators $\{\hat{a}_j, \hat{a}_j^\dagger\}$ and the (abstract) position and momentum operators $\{\hat{x}_j, \hat{p}_j\}$ of the individual field modes, we introduce canonical position and momentum variables [66, 28]

$$Q_n = \frac{1}{\sqrt{2\epsilon_n}}(\alpha_n + \alpha_n^*) \quad (3.10)$$

$$P_n = i\sqrt{\frac{\epsilon_n}{2}}(\alpha_n^* - \alpha_n). \quad (3.11)$$

The usual Hamilton's equations [117, 252] for Q_n and P_n then yield $id\alpha_n/dt = \partial H/\partial \alpha_n^*$, and the complex-conjugate expression. Introducing the *projected functional derivative* operator [104, 28]

$$\frac{\bar{\delta}}{\bar{\delta}\psi(\mathbf{x})} = \sum_{n \in \mathbf{C}} \phi_n^*(\mathbf{x}) \frac{\partial}{\partial \alpha_n}, \quad (3.12)$$

and its conjugate, the evolution equations for the classical-field amplitudes can be summarised by the functional equation

$$i\hbar \frac{\partial \psi(\mathbf{x})}{\partial t} = \frac{\bar{\delta} H_{\text{CF}}}{\bar{\delta} \psi^*(\mathbf{x})}, \quad (3.13)$$

for the projected classical field. The resulting equation of motion

$$i\hbar \frac{\partial \psi(\mathbf{x})}{\partial t} = \mathcal{P} \left\{ (H_{\text{sp}} + U_0 |\psi(\mathbf{x})|^2) \psi(\mathbf{x}) \right\}, \quad (3.14)$$

is called the *projected* Gross-Pitaevskii equation (PGPE) (see, e.g., [28]).

Conservation laws and Ehrenfest relations

We now remark on some conservation laws satisfied by the evolution of the PGPE. In general, there will be multiple *first integrals* of the PGPE evolution. The first integrals are functionals of the field which are conserved by the Hamiltonian evolution. In general applications of the PGPE, there are always at least two such functionals: the field energy

Chapter 3. Classical-field theory

$E[\psi]$ and normalisation $N[\psi]$. More generally, there may be other conserved quantities which are conjugate to symmetries of the classical-field Hamiltonian (3.8) (e.g., conserved angular momentum in the case of a rotationally invariant single-particle Hamiltonian).

Here we discuss the appearance of the normalisation and energy first integrals in the PGPE formalism. First, we note that the Hamiltonian (3.8) is invariant under (global) phase rotations $\psi(\mathbf{x}) \rightarrow \psi(\mathbf{x})e^{i\theta}$. The group of such transformations is generated by the classical-field normalisation $N = \int d\mathbf{x} |\psi(\mathbf{x})|^2$, which is therefore conserved [117]. Introducing the projected functional Poisson bracket [28]

$$\{F, G\} = \int d\mathbf{x} \left[\frac{\bar{\delta}F}{\bar{\delta}\psi(\mathbf{x})} \frac{\bar{\delta}G}{\bar{\delta}\psi^*(\mathbf{x})} - \frac{\bar{\delta}F}{\bar{\delta}\psi^*(\mathbf{x})} \frac{\bar{\delta}G}{\bar{\delta}\psi(\mathbf{x})} \right], \quad (3.15)$$

it can be shown explicitly that $\{N, H_{\text{CF}}\} = 0$. This conservation of the field normalisation as a result of the invariance of the Hamiltonian H_{CF} under such phase transformations is the classical precursor to the ‘gauge’ symmetry of the quantised field discussed in section 2.2.2. Furthermore, if the potential $\delta V(\mathbf{x})$ is time-independent it follows immediately [117, 252] that the energy (3.8) of the classical field is conserved by the PGPE evolution.

We can also derive *Ehrenfest* relations for the evolution of ‘expectation values’ of single-particle operators, obtained by averaging over the (instantaneous) classical-field configuration

$$\bar{A} \equiv \int d\mathbf{x} \psi^*(\mathbf{x}) A \psi(\mathbf{x}). \quad (3.16)$$

The standard (non-projected) Gross-Pitaevskii equation satisfies the conventional Ehrenfest relations for the quantities of immediate interest: the position \mathbf{x} , momentum \mathbf{p} and angular momentum \mathbf{L} (see, e.g., [44]). The corresponding results for the PGPE also contain corrections due to the projection operation, which are most straightforwardly derived by rewriting the PGPE as [36, 34]

$$i\hbar \frac{\partial \psi}{\partial t} = (1 - \mathcal{Q}) \left\{ (H_{\text{sp}} + U_0 |\psi(\mathbf{x})|^2) \psi(\mathbf{x}) \right\}, \quad (3.17)$$

where we have introduced the projector \mathcal{Q} which is the orthogonal complement of \mathcal{P} , i.e.

$$\mathcal{Q}f(\mathbf{x}) \equiv (1 - \mathcal{P})f(\mathbf{x}) = \sum_{n \in \mathbf{I}} \phi_n(\mathbf{x}) \int d\mathbf{x}' \phi_n^*(\mathbf{x}') f(\mathbf{x}'). \quad (3.18)$$

We then find the generalised Ehrenfest relations

$$\frac{d\bar{\mathbf{x}}}{dt} = \frac{\bar{\mathbf{p}}}{m} + Q_{\mathbf{x}} \quad (3.19)$$

$$\frac{d\overline{\mathbf{p}}}{dt} = -\overline{\nabla V(\mathbf{x}, t)} + Q_{\mathbf{p}} \quad (3.20)$$

$$\frac{d\overline{\mathbf{L}}}{dt} = -\frac{i}{\hbar}\overline{\mathbf{L}V(\mathbf{x}, t)} + Q_{\mathbf{L}} \quad (3.21)$$

$$\frac{dH_{\text{CF}}}{dt} = \overline{\frac{\partial V(\mathbf{x}, t)}{\partial t}}, \quad (3.22)$$

where the *boundary terms* are

$$Q_A = \frac{2}{\hbar} \text{Im} \left[(\delta V(\mathbf{x}, t) + U_0 |\psi(\mathbf{x})|^2) \mathcal{Q} \{A\psi(\mathbf{x})\} \right]. \quad (3.23)$$

It is clear from (3.23) that the boundary terms vanish when $[\mathcal{Q}, A] \equiv -[\mathcal{P}, A] = 0$, in which case we regain the Ehrenfest relations of the continuum field theory. This result will be of particular importance when we apply the PGPE to rotating systems in later chapters.

3.3 Ergodicity

The PGPE (3.14) is, fundamentally, an equation of motion which describes the *dynamics* of the classical field $\psi(\mathbf{x}, t)$. However, an important feature of the dynamical system described by the Hamiltonian (3.8) is that it is *ergodic* [160], which is to say that trajectories of the PGPE sample the microcanonical density of the field densely. For given values $\{I_i\}$ of the first integrals $\{F_i\}$ of the PGPE system, the microcanonical density is the probability distribution over field configurations which has a constant nonzero value on the sub-manifold defined by $F_i = I_i$, and is zero elsewhere. In the simple case in which $E[\psi]$ and $N[\psi]$ are the only first integrals, with values E_0 and N_0 respectively, the microcanonical density is

$$P_{\mu}[\psi; E_0; N_0] = \begin{cases} \text{const} & E[\psi] = E_0, N[\psi] = N_0 \\ 0 & \text{else,} \end{cases} \quad (3.24)$$

where the constant is chosen such that the integral of the density over all field configurations is unity. In principle, all equilibrium properties of the microcanonical classical-field system are given by averages in the density (3.24). As the evolution is ergodic, almost any initial condition with $E[\psi] = E_0$ and $N[\psi] = N_0$ will evolve to a configuration (microstate) which is *typical* of the appropriate thermal equilibrium of the field⁴. After this

⁴We use the phrase ‘almost any’ to signify that there are some field configurations (a set of measure zero [232]) which will not evolve in an ergodic manner. An important case is initial conditions which exhibit a symmetry which is preserved by the Hamiltonian evolution.

equilibration of the field, the field essentially samples from the density (3.24) uniformly over time. This provides a powerful methodology for calculating the equilibrium properties of the classical field: we simply allow the system to attain its own equilibrium, and then *observe* its equilibrium behaviour directly.

It is worth elaborating somewhat on the nature of this equilibration, as the way this apparently irreversible behaviour arises from the purely deterministic equation (3.14) is a potential source of confusion. An initial field configuration that we choose may appear far from the typical configurations of the field we observe in equilibrium. However, in equilibrium the system will eventually fluctuate so that all such unlikely configurations occur, as their probability in the microcanonical density is of course the same as any other microstate consistent with $P_\mu[\psi]$ ⁵. The subsequent ‘relaxation’ we observe is simply the system continuing with its evolution after such a fluctuation, and is formally reversible. In particular, the information describing the initial state is preserved in the precise details of the fluctuations of the field once it has reached equilibrium [57], and the initial state can be regained by reversing the evolution of the system. However, in practical applications this reversibility is inevitably lost after some time due to the buildup of numerical error [160].

3.3.1 Correlation functions

A particularly useful application of the ergodicity of the PGPE system is the calculation of *correlation functions* of the classical field, which take the form of averages in the microcanonical density $P_\mu[\psi]$ (which we denote by $\langle \cdots \rangle_\mu$) of a functional $F[\psi(\mathbf{x})]$ of the classical field. Such averages provide various characterisations of the properties of the classical field in equilibrium, and are analogous to the quantum-statistical-mechanical expectations of field variables obtained by tracing over the quantum density matrix $\langle F[\hat{\psi}(\mathbf{x})] \rangle = \text{Tr}\{\hat{\rho}F[\hat{\psi}(\mathbf{x})]\}$. The (hypothesised) ergodicity is formalised by the identification of the microcanonical average $P_\mu[\psi]$ with the infinite-time limit of an average along any given trajectory of the field [160]:

$$\langle F[\psi(\mathbf{x})] \rangle_\mu = \lim_{\theta \rightarrow \infty} \left\{ \frac{1}{\theta} \int_0^\theta dt F[\psi(\mathbf{x}, t)] \right\}. \quad (3.25)$$

This is a very useful identity, as identifying the full set of classical-field configurations which satisfy the first-integral constraints is almost certainly impossible, due to the non-linearity of the classical-field energy functional (3.8) [28]. In practice, we average the

⁵They are *unlikely* configurations, however, in the sense that their number of such configurations (or more accurately the volume of microcanonical density that they occupy) is comparatively small.

functional $F[\psi(\mathbf{x}, t)]$ over a finite period of field evolution. The trajectories of the PGPE thus provide for a kind of *Monte Carlo* sampling of the microcanonical density of the field, with the distinction that the ‘randomness’ of the sampling arises from the ergodicity of the deterministic trajectories, rather than any explicit stochastic element. Intuitively, we expect the results of the averaging to be accurate estimates of the corresponding microcanonical averages, provided that we average over a period which is long compared to the slowest time scale in the problem (at equilibrium), which is of the order of the slowest harmonic trapping period.

Field covariance

A correlation function of the classical field of particular interest is the *covariance matrix*

$$G(\mathbf{x}, \mathbf{x}') = \langle \psi^*(\mathbf{x}) \psi(\mathbf{x}') \rangle_\mu. \quad (3.26)$$

This quantity is the obvious classical-field analogue of the one-body density matrix $\rho_1(\mathbf{x}, \mathbf{x}')$ introduced in section 2.2.1⁶. In particular, $G(\mathbf{x}, \mathbf{x}')$ is, like $\rho_1(\mathbf{x}, \mathbf{x}')$, Hermitian, and can be written in the diagonal form

$$G(\mathbf{x}, \mathbf{x}') = \sum_i n_i \chi_i^*(\mathbf{x}) \chi_i(\mathbf{x}'), \quad (3.27)$$

where the real constants $\{n_j\}$, indexed in order of decreasing magnitude, are the occupations of the corresponding modes $\{\chi_j(\mathbf{x})\}$. By analogy to the Penrose-Onsager [201] definition of condensation in the quantum-many-body system (section 2.2.1), condensation in the classical field is signaled by the most highly occupied mode $\chi_0(\mathbf{x})$ having an occupation n_0 which is significantly larger than all other occupations n_i . This definition in terms of correlations in the microcanonical density is an unambiguous measure of condensation in the simple equilibrium regimes in which it is applicable.

Beyond the microcanonical density

A major theme of this thesis is the generalisation of the notion of correlation functions of the field to situations in which the correlations of interest are in some way *concealed* in the formal microcanonical density itself. For example, the microcanonical density of the field inherits the U(1) phase symmetry of the classical-field Hamiltonian (section 3.2.2). However, we show in chapter 5 that over comparatively short time periods, the classical

⁶In light of this analogy, in the classical-field context we will sometimes simply refer to the covariance matrix $G(\mathbf{x}, \mathbf{x}')$ as the one-body density matrix.

field effectively samples a modified ensemble which breaks the $U(1)$ phase symmetry of the Hamiltonian (3.8), and so by considering the fluctuation statistics of the field over such short periods we can obtain nonzero values for *symmetry-broken* correlation functions, in analogy to the symmetry-broken approach to the second-quantised field discussed in section 2.2.2. Similarly in chapter 6, we consider a scenario in which the classical-field condensate breaks the *rotational* symmetry of the classical-field Hamiltonian. By considering the fluctuations of the field on short time scales we can extract the symmetry-broken condensate orbital, which is concealed in the formal rotationally-invariant microcanonical density. Another scenario which necessitates considering field correlations beyond the microcanonical density is the case of *nonequilibrium* fields, which we consider in chapters 7 and 9. In those scenarios, we are interested in the *relaxation* of the classical field to a (new) thermal equilibrium, and we thus require means of characterising the field during this relaxation period. We will find in such cases that the separation of time scales between the slow ‘macroscopic’ relaxation of the field and the rapid motion of thermal fluctuations of the field is such that we can obtain meaningful characterisations of the field from averages taken over intermediate time scales.

3.3.2 Thermodynamics

Microcanonical thermodynamics

The second important application of the ergodicity of the field is to the calculation of the thermodynamic parameters of the system: the temperature and chemical potential. As the PGPE system is Hamiltonian and ergodic, its thermodynamics are in principle described completely by the microcanonical *entropy function*

$$S = k_B \ln \int d\mathbf{\Gamma} \delta(H(\mathbf{\Gamma}) - E) \prod_i \delta(F_i(\mathbf{\Gamma}) - I_i), \quad (3.28)$$

where $H(\mathbf{\Gamma})$ is the classical-field Hamiltonian (3.8) written as a function of the phase-space variables $\mathbf{\Gamma}$, and the F_i are the other conserved first integrals of the Hamiltonian evolution (e.g., the field normalisation, see section 3.2.2). The equilibrium temperature of the microcanonical system is formally given by

$$\frac{1}{k_B T} = \left(\frac{\partial S}{\partial E} \right)_{F_i}, \quad (3.29)$$

where the subscript F_i indicates that the derivative is taken with all other constants of motion held fixed. Calculating the temperature using this definition however requires the

calculation of the entropy function as a function of energy, which is again an intractable task. However, a comparatively recent advance in statistical mechanics was provided by Rugh [224], who showed that the temperature can instead be expressed as an average in the microcanonical density

$$\frac{1}{T(E)} = \langle \mathcal{F}[H] \rangle_\mu, \quad (3.30)$$

where $\mathcal{F}[H]$ is a given functional of the system Hamiltonian [224]. This provides a great advantage as it allows us to calculate the temperature by an ergodic averaging procedure. However, the evaluation of the functional \mathcal{F} in terms of the classical-field canonical variables $\{Q_j, P_j\}$ is very complicated, and so this technology has only been implemented for the cases of plane-wave [66] and Hermite-basis projection [64]. In this thesis we consider some scenarios with an additional conserved first integral (angular momentum) for which the Rugh approach would require significant extension, and we also consider nonequilibrium fields, for which the Rugh approach is not applicable. We therefore need also to consider other approaches to estimating the temperature, based on the principle of equipartition of energy in the microcanonical system.

Equipartition and the classical-field distribution function

Although the thermodynamics of the interacting classical field system can only be evaluated numerically, we can gain some insight into the thermodynamic behaviour of the classical-field system by considering the noninteracting limit. In this limit, the Hamiltonian becomes quadratic in the field variables $\{\alpha_n, \alpha_n^*\}$, and so the principle of *equipartition* of energy applies. It is most convenient to consider the situation in the grand-canonical ensemble. The equipartition theorem [197] then implies that a single-particle mode with energy ϵ_k has mean occupation

$$n_k = \frac{k_B T}{\epsilon_k - \mu}. \quad (3.31)$$

A crucial point is that this distribution is obtained from the formal Bose-Einstein distribution (2.10) in the limit of high occupations $n_k \gg 1$, which occurs when the parameter $(\epsilon_k - \mu)/k_B T \ll 1$. Formally, the equipartition distribution is obtained as the first term in the expansion of the Bose distribution (2.10) in this parameter [30]. As is well known, the distribution (3.31), when applied to a continuum field theory (with an infinite number of modes) leads to an ultraviolet catastrophe: each mode of the field contains equal (finite) energy, leading to an infinite total energy. Of course, in the classical field theory with a finite field energy, we face the inverse problem [28]: as the cutoff is raised, the energy in the field will be divided over an increasing number of modes, and the occupation of

each individual mode tends to zero as the cutoff is raised to infinity⁷. In the second-quantised field theory, the occupation of each mode is quantised, and so at high energies where $n_k \ll 1$ the decay becomes exponential in the energy (approaching a Boltzmann distribution), regularising this divergent behaviour.

The distribution (3.31) provides a useful guide to the thermodynamic behaviour of the classical-field system, and allows us to understand how a classical-field model can be used to emulate the statistical mechanics of a Bose-field system. The central point is that the distribution (3.31) shares the *saturation* property of the Bose distribution. For a given temperature, the total excited-state population is bounded from above, as for the Bose distribution, except that the bound in this case depends crucially on the fact that there are only a finite number of excited states below the cutoff. As a result, the classical-field distribution describes condensation in the lowest-energy mode, with a critical temperature depending explicitly on the cutoff energy E_R . In appendix A we derive the general (geometry-independent) form of the transition temperature. For the most experimentally relevant case of three-dimensional harmonic trapping (with density of states $g(\epsilon) = \epsilon^2/2(\hbar\bar{\omega})^3$, where $\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}$ is the geometric mean trapping frequency) we obtain

$$k_B T = \frac{N}{(\hbar\bar{\omega})^3 E_R^2}. \quad (3.32)$$

Going beyond the noninteracting case, we expect the equipartition principle to apply to the noncondensed modes of the classical field, to the extent to which they can be considered to be weakly interacting. Lobo *et al.* [170] have used this principle to estimate the temperature of a classical field, by assuming that the energy of the field above that of the corresponding ground state is equipartitioned over $\mathcal{M} - 1$ noninteracting modes. In later chapters we take a related approach, and construct a spatial fitting function based on the assumption of equipartition over the thermal modes of the field, and a Hartree-Fock model for the field interactions. Fitting this function to the spatial distribution of the field allows us to estimate effective thermodynamic parameters of the field in nonequilibrium scenarios.

3.4 Truncated-Wigner approximation

In the discussion above, we have obtained a classical-field equation of motion simply by replacing the quantised field operator $\hat{\psi}$ with a classical-valued analogue ψ . Here we

⁷In practice this should not occur for the Hamiltonian field dynamics, as the perturbing effect of field interactions and thus the coupling to higher modes diminishes rapidly for increasing mode energies, so the ‘creep’ of population to higher energies will eventually terminate.

discuss how the same equation of motion can be obtained by more formal means. In this *truncated-Wigner* approach, the classical equation of motion (3.14) is obtained as an approximate equation of motion for classical trajectories which sample the quantum evolution of the system. In this approach, formal quantum expectation values can be obtained through simple correspondences with averages over an ensemble of the classical trajectories. In this thesis we consider scenarios of strongly thermal dynamics, and the relaxation of atomic fields over long time scales, beyond the strict validity of the formal truncated-Wigner method. Nevertheless, considering the truncated-Wigner approximation lends us further insight into the content of the classical-field method, and in particular shows us how to avoid some potential ambiguities and pathologies in our classical-field simulations.

3.4.1 The Wigner distribution

Phase-space methods

An important tool for the study of quantum-mechanical systems developed and frequently employed in quantum optics research is the concept of a *quantum-classical correspondence*. The basic idea [106] is to establish a correspondence between the quantum density matrix $\hat{\rho}$ and a *classical function* $P(\mathbf{z})$ of *phase-space* variables (e.g., x and p), which we will denote generically by a vector \mathbf{z} . The function $P(\mathbf{z})$ is a kind of probabilistic distribution over these phase-space variables, however, in order to completely describe the quantum state of the system, the function $P(\mathbf{z})$ generally fails to exhibit all the familiar properties of a classical probability distribution; e.g., $P(\mathbf{z})$ may be non-positive-definite, or even singular. An important feature of this approach is that the *evolution equation* for $\hat{\rho}$ can be transformed into a classical partial differential equation for $P(\mathbf{z})$. This allows us to apply many techniques of classical statistical mechanics to the analysis and solution of the evolution equation for $P(\mathbf{z})$ [46].

A particular advantage of the corresponding classical formulation of a quantum-dynamical system arises in scenarios in which the classical equation of motion for $P(\mathbf{z})$ is (or can be approximated by) a *Fokker-Planck equation*

$$\begin{aligned} \frac{dP(\mathbf{z}, t)}{dt} = & - \sum_j \frac{\partial}{\partial z_j} [A_j(\mathbf{z}, t) P(\mathbf{z}, t)] \\ & + \frac{1}{2} \sum_{jk} \frac{\partial^2}{\partial z_j \partial z_k} \left\{ [\bar{\mathbf{B}}(\mathbf{z}, t) \bar{\mathbf{B}}^T(\mathbf{z}, t)]_{jk} P(\mathbf{z}, t) \right\}, \end{aligned} \quad (3.33)$$

which we have expressed in terms of the *drift vector* $\mathbf{A}(\mathbf{z}, t)$ and *diffusion matrix* $\bar{\mathbf{B}}(\mathbf{z}, t)$.

In this case we can exploit the correspondence between such a Fokker-Planck equation and the (Ito) *stochastic differential equation* (SDE)

$$d\mathbf{z} = \mathbf{A}(\mathbf{z}, t)dt + \bar{\mathbf{B}}(\mathbf{z}, t)d\mathbf{W}(t), \quad (3.34)$$

where $d\mathbf{W}(t)$ is a vector of independent Wiener processes [101]. The SDE (3.34) includes the effects of fluctuations such that an ensemble of independent stochastic solutions $\{\mathbf{z}_i(t)\}$ exhibits the statistics described by the solution of the Fokker-Planck equation (3.33)⁸. The primary advantage of this formulation is the comparative numerical simplicity of integrating such a stochastic equation, as opposed to integrating the Fokker-Planck equation for $P(\mathbf{z}, t)$ itself. Moreover, as individual stochastic trajectories (corresponding to distinct realisations of the noise process $d\mathbf{W}(t)$) are such that their *ensemble* statistics correspond to the statistics of the quantum-mechanical ensemble, individual trajectories can be interpreted (with caution) as representative of particular observed trajectories of the true quantum system [28].

The Wigner distribution

We consider here a particular quasiprobability distribution which is relevant to the classical-field approximation used in this thesis. It is in fact historically the first such distribution, introduced by Wigner in 1932 [263]. Its meaning is best understood by first introducing it for a single-particle system, in terms of the fundamental dynamical variables x and p . We define

$$W(x, p) = \frac{1}{\pi\hbar} \int dy \langle x - y | \hat{\rho} | x + y \rangle e^{2ipy/\hbar}, \quad (3.35)$$

which is a kind of phase-space density in x and p , analogous to that in Hamiltonian mechanics and statistical mechanics [252]. In particular, we find for the marginal probabilities $\int dp W(x, p) = \langle x | \hat{\rho} | x \rangle$ and $\int dx W(x, p) = \langle p | \hat{\rho} | p \rangle$ [106]. However, unlike a classical probability density, $W(x, p)$ can legitimately acquire a negative value at a given point in phase space⁹. Non-positivity of the Wigner function is associated with quantum interference and other nonclassical effects. Regarded as a probability distribution, it can be shown that the moments of the Wigner distribution correspond to the quantum expectation values

⁸The relation between the Fokker-Planck and SDE formulations is thus closely analogous to that between the Schrödinger and Heisenberg pictures of classical and quantum mechanics [252, 101].

⁹The interpretation of quantum mechanics as a theory which admits such negative probabilities is discussed, for example, by Feynman [91].

of *symmetrised* products of the variables x and p , i.e.

$$\langle \{x^m p^n\}_{\text{sym}} \rangle = \int dx \int dp W(x, p) x^m p^n. \quad (3.36)$$

In the second-quantised description of electromagnetic and atomic fields, each mode of the field is dynamically equivalent to a harmonic oscillator, and so in quantum and atom optics applications, it is most convenient to define the Wigner function in terms of the ladder operators \hat{a} and \hat{a}^\dagger of these oscillators. This representation is straightforwardly related to the (x, p) representation [129], due to the simple canonical relation between the two representations (cf. equations (3.10) and (3.11)). We find for the Wigner function of a single mode of the field

$$W(\alpha, \alpha^*) = \frac{1}{\pi^2} \int d^2\lambda e^{-\lambda\alpha^* + \lambda^*\alpha} \text{Tr}\{\hat{\rho} e^{\lambda\hat{a}^\dagger - \lambda^*\hat{a}}\}, \quad (3.37)$$

and the relation of moments of this function to quantum expectation values is

$$\langle \{(a^\dagger)^m a^n\}_{\text{sym}} \rangle = \int d^2\alpha (\alpha^*)^m \alpha^n W(\alpha, \alpha^*) \equiv \langle (\alpha^*)^m \alpha^n \rangle_W. \quad (3.38)$$

Furthermore we can deduce correspondences between the *pre*– or *post*– multiplication of the density matrix by a ladder operator and the action of a *classical* operator on the Wigner function. We find [106]

$$\hat{a}\hat{\rho} \leftrightarrow \left(\alpha + \frac{1}{2} \frac{\partial}{\partial \alpha^*} \right) W(\alpha, \alpha^*) \quad (3.39)$$

$$\hat{a}^\dagger \hat{\rho} \leftrightarrow \left(\alpha^* - \frac{1}{2} \frac{\partial}{\partial \alpha} \right) W(\alpha, \alpha^*) \quad (3.40)$$

$$\hat{\rho}\hat{a} \leftrightarrow \left(\alpha - \frac{1}{2} \frac{\partial}{\partial \alpha^*} \right) W(\alpha, \alpha^*) \quad (3.41)$$

$$\hat{\rho}\hat{a}^\dagger \leftrightarrow \left(\alpha^* + \frac{1}{2} \frac{\partial}{\partial \alpha} \right) W(\alpha, \alpha^*). \quad (3.42)$$

These correspondences allow us to convert a quantum (von Neumann or master) equation for the density matrix into a classical partial-differential equation for the Wigner function.

3.4.2 Truncated-Wigner approximation for the interacting Bose gas

Equations of motion

We now discuss the equations of motion that result from the application of the Wigner formalism to the low-energy effective Hamiltonian (3.6) of the dilute Bose gas. We begin with the Schrödinger-picture representation of the quantum dynamics, where the density-matrix $\hat{\rho}$ evolves according to the von Neumann equation

$$\begin{aligned} i\hbar \frac{\partial \hat{\rho}}{\partial t} &= [\hat{H}, \hat{\rho}] \\ &= \int d\mathbf{x} \left\{ \hat{\psi}^\dagger H_{\text{sp}} \hat{\psi} \hat{\rho} - \hat{\rho} \hat{\psi}^\dagger H_{\text{sp}} \hat{\psi} + \frac{U_0}{2} \left(\hat{\psi}^\dagger \hat{\psi}^\dagger \hat{\psi} \hat{\psi} \hat{\rho} - \hat{\rho} \hat{\psi}^\dagger \hat{\psi}^\dagger \hat{\psi} \hat{\psi} \right) \right\}. \end{aligned} \quad (3.43)$$

As the quantum system contains a finite number of modes (those spanning the low-energy region \mathbf{C}), the Wigner function (3.37) must be generalised to the multimode case. Fortunately this generalisation is straightforward, and is discussed (for e.g.) in [28]. For notational compactness we express the multimode Wigner function as a *functional* distribution $W(\{\psi, \psi^*\})$; using the projected functional differentiation operators (3.12), the (projected) functional operator correspondences are obtained from (3.39) by replacing $\hat{a} \rightarrow \hat{\psi}$, $\alpha \rightarrow \psi$ and $\partial/\partial\alpha \rightarrow \bar{\delta}/\bar{\delta}\psi$, *etc.* We thus find the Wigner-function evolution equation [246]

$$\begin{aligned} \frac{\partial W(\{\psi, \psi^*\})}{\partial t} &= \frac{i}{\hbar} \int d\mathbf{x} \left\{ \frac{\bar{\delta}}{\bar{\delta}\psi} \left[H_{\text{sp}} + U_0(|\psi|^2 - \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x})) \right] \psi \right. \\ &\quad \left. - \frac{\bar{\delta}}{\bar{\delta}\psi^*} \left[H_{\text{sp}} + U_0(|\psi|^2 - \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x})) \right] \psi^* \right. \\ &\quad \left. - \frac{U_0}{4} \left(\frac{\bar{\delta}^3}{\bar{\delta}\psi^2 \bar{\delta}\psi^*} \psi - \frac{\bar{\delta}^3}{\bar{\delta}\psi \bar{\delta}\psi^{*2}} \psi^* \right) \right\} W(\{\psi, \psi^*\}). \end{aligned} \quad (3.44)$$

We note that there are no second-derivative terms in this evolution equation. However, *third-order* derivative terms appear, which correspond to the effects of dynamical quantum fluctuations. These quantum effects are expected to be negligible for modes with high occupations, as discussed in section 3.1. Neglecting these terms we thus obtain an equation for $W(\{\psi, \psi^*\}, t)$ which describes *nondiffusive* (classical Liouvillian) evolution of the Wigner function. Consequently, the corresponding ‘stochastic’ differential equation which samples the evolution of this distribution (see equation (3.34)) is in fact the

deterministic differential equation

$$i\hbar \frac{\partial \psi(\mathbf{x}, t)}{\partial t} = H_{\text{sp}} \psi(\mathbf{x}) + \mathcal{P} \left\{ U_0 (|\psi(\mathbf{x})|^2 - \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x})) \psi(\mathbf{x}) \right\}. \quad (3.45)$$

The truncated-Wigner approximation is valid provided that the dynamical quantum corrections encoded in the neglected third-derivative terms play little role in the evolution of the underlying physical system. An unambiguous validity criterion for the truncated-Wigner method is thus that all modes in the region \mathbf{C} are highly occupied. Indeed in the ‘classical’ limit $U_0 \rightarrow 0$, $N[\psi] \rightarrow \infty$, $U_0 N[\psi] = \text{const}$, the third-derivative terms in equation (3.44) vanish and the equation of motion (3.45) is thus the asymptotically exact equation of motion for the system in this limit [28]. In this limit, the projected delta function which appears in equation (3.45) as an additional (external) potential term also vanishes, and so this term is typically neglected [190, 34], and we regain the PGPE (equation (3.14)). In fact, Norrie *et al.* [190] have derived a weaker condition for the validity of the evolution in terms of the local density of the system in position space. However, the *formal* interpretation of the stochastic trajectories of equation (3.45) as samples of the Wigner distribution evolution is not valid on long time scales, due to the ergodic thermalisation of the field trajectories, as we discuss in section 3.4.3.

Initial conditions

Although the evolution equation (3.45) is itself deterministic, a stochastic element (corresponding to quantum fluctuations) does enter into the truncated Wigner description through the *initial conditions* of the classical-field trajectories. Generalising the relation (3.38) to the multimode case [28] we find for example

$$\begin{aligned} \langle \psi^*(\mathbf{x}) \psi(\mathbf{x}') \rangle_{\text{W}} &\equiv \int \bar{\mathcal{D}}^2 \psi \psi^*(\mathbf{x}) \psi(\mathbf{x}') W(\{\psi, \psi^*\}) \\ &= \langle \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}') \rangle + \frac{1}{2} \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x}'), \end{aligned} \quad (3.46)$$

where we have introduced the projected functional integration operator $\int \bar{\mathcal{D}}^2 \psi = \prod_{n \in \mathbf{C}} \int d^2 \alpha_n$. Thus even in the absence of any population ($\langle \hat{\psi}^\dagger \hat{\psi} \rangle = 0$), the density of the field in the Wigner distribution is finite ($\langle |\psi(\mathbf{x})|^2 \rangle_{\text{W}} = (1/2) \delta_{\mathcal{P}}(\mathbf{x}, \mathbf{x})$). This density corresponds to the *vacuum fluctuations* of the field, i.e., the zero-point occupation of each mode of the field¹⁰. For noninteracting modes, this statistical spread in the

¹⁰In the limit that the cutoff height is raised to infinity, the projected delta function becomes a true delta function, and we regain the well-known result that there is infinite population (and thus energy) contained in the zero-point occupation of a (continuum) bosonic field.

Wigner function is included simply by populating the mode amplitudes α_k according to $\langle \alpha_k^* \alpha_l \rangle = (n_k + 1/2) \delta_{kl}$, with $\langle \alpha_k \rangle = \langle \alpha_k^* \alpha_l \rangle = 0$, where n_k is the (real) thermal population. This can easily be achieved by, for example, sampling the $\{\alpha_j\}$ from an appropriate Gaussian distribution. In the case of the interacting fields such as we consider in this thesis, a single mode basis is not an appropriate basis for adding noise and, in general, we can only attempt to populate the field with noise in a way that best approximates the true many-body state we wish to represent. There are many different approaches to constructing the noise, depending on the system and state being modelled [246, 102, 239, 212, 34, 28], as we discuss in chapter 9.

3.4.3 Relevance to the PGPE formalism

Perhaps the greatest importance of the truncated-Wigner derivation of the projected Gross-Pitaevskii equation is that it shows that the classical-field equation of motion can be obtained from the full many-body theory in a *controlled* approximation. In fact, Polkovnikov [210, 211] has shown how the full quantum dynamics can be formally expanded about a Gross-Pitaevskii equation-like description, and that the truncated-Wigner method appears as the inclusion of quantum corrections to leading order in this expansion.

Another important result obtained from the truncated-Wigner formalism is the appearance of the initial noise associated with quantum fluctuations. In general, this noise allows for the inclusion of spontaneous processes in the GP-type description, which otherwise describes only coherent processes. This also allows us to avoid some pathologies of a pure-GP description of the field. In fact, the symmetries of the pure-GP condensate mode can be such that the deterministic evolution of the PGPE is constrained from thermal equilibration. This can be equivalently viewed as a state which is unstable to decay by spontaneous processes being unable to do so in the absence of added fluctuations, or as the initial GP state lying on the ‘set of measure zero’ of microcanonical trajectories which do not exhibit ergodic behaviour. The truncated-Wigner removes this issue unambiguously, by showing that for a classical field to be sensibly interpreted as a beyond-mean-field description of the quantum Bose field, it must include an irreducible level of fluctuations, which *break* all symmetries of the naive GP state, and seed all spontaneous transitions of the field. This aspect will be important in chapter 9 when we consider the evolution of an initially zero-temperature condensate following a *dynamical instability*, in which unstable excitations of the condensate mode grow exponentially and spontaneously.

It is important to note that the truncated-Wigner method is only an approximation, and in particular, the error associated with the neglect of the cubic-noise terms grows with time during the system evolution. Indeed we have already seen this from our consideration

of the PGPE trajectories: as the time $t \rightarrow \infty$, the Hamiltonian trajectories come to an ergodic equilibrium. Thus the moments of the Wigner function at long times correspond to moments of the equilibrium microcanonical density of the classical field, and even in the high-occupation regime, the formal operator expectations $\langle \hat{a}_k^\dagger \hat{a}_k \rangle$ we obtain from the Wigner ensemble will be distorted by the subtraction of the half quanta of ‘noise’ from each mode of the classical-field distribution. We thus face again the fundamental issue that the neglect of the Bose commutation relations in the evolution of the field leads to the loss of the exact Bose statistics of the field. A potentially problematic consequence of this behaviour is that the energy and normalisation added to the field as a representation of the vacuum noise leads in general to a spurious raising of the temperature of the field. This consequently affects damping rates in the field, which are strongly dependent on the temperature [82]. Sinatra *et al.* [239] thus suggested the additional condition $T_{\text{class}} - T \ll T$ on the temperature of the field, in relation to the temperature T_{class} of the Hamiltonian field at microcanonical equilibrium. However, this condition does not necessarily apply in strongly nonequilibrium scenarios in which the temperature changes dramatically, and in chapter 9 we consider the rethermalisation of a classical field following a dynamical instability at *zero temperature*, but nevertheless find that the energy (and angular momentum) added to the system in the form of vacuum noise is small compared to that ‘liberated’ from the mean field in the transition to the new equilibrium.

The stochastic Gross-Pitaevskii equation

Here we briefly discuss the *stochastic* Gross-Pitaevskii equation (SGPE) developed by Gardiner and coworkers [102, 104]. This formalism is closely related to the finite-temperature projected Gross-Pitaevskii equation, and the formal truncated-Wigner method. The essence of the treatment is that an *open systems* approach is taken to deriving an equation of motion for the condensate-band density operator, including the dissipative effect of its coupling to the above-cutoff region. The above-cutoff atoms are assumed to be at thermal equilibrium, and form a *thermal bath* for the below cutoff atoms. The evolution equation for the below-cutoff atoms is then a *master equation*, derived using a high-temperature expansion. Mapping this master equation onto an evolution equation for the Wigner function we find proper diffusive terms (i.e., second order derivatives). The appropriate stochastic equations for the classical field then include explicit noise and damping terms, resulting from the coupling to the thermal bath. The coupling terms include the exchange of atoms between the condensate band and the bath, and the stochastic Gross-Pitaevskii equation is thus a *grand-canonical* method for the Bose field. To ensure the validity of the approximations made in the derivation, it is limited to the study of comparatively

Chapter 3. Classical-field theory

high-temperature Bose fields ($T \gtrsim 0.7T_c$), but allows in particular for the study of the response of the Bose field to *quenches* of the thermal parameters of the bath [37, 260].

Chapter 4

Numerical methods

In this chapter, we outline the numerical integration of the projected Gross-Pitaevskii equation (PGPE). We discuss the reformulation of the PGPE in terms of dimensionless quantities, interaction-picture methods for the temporal integration of the PGPE, the spectral-Galerkin approach to its spatial integration, and Gaussian quadrature rules for the efficient evaluation of the integrals that arise in this formulation. We then contrast the PGPE approach with other numerical discretisations of the Gross-Pitaevskii theory, and explain why the projection operation is necessary for the accurate simulation of the classical-field dynamics.

4.1 Numerical integration of the projected Gross-Pitaevskii equation

4.1.1 Dimensionless formalism

In implementing the numerical solution of the projected Gross-Pitaevskii equation, we reformulate it in terms of *dimensionless quantities*. This serves two purposes: first, it ensures that the numerical quantities manipulated during the integration are of reasonable magnitude, i.e., we avoid carrying around factors of very large or (as is typically the case in describing quantum-mechanical systems in SI units) very small numbers that can lead to numerical round-off error. The second reason is that by identifying appropriate ‘natural’ units of the system we often find that the behaviour of the system depends on an irreducible set of parameters which is smaller than the set of physical (dimensioned) parameters. The systems we treat in this thesis are harmonically trapped, and this trapping provides for obvious definitions for natural units. In general, we consider systems which are oblate, and so we base our units on the in-plane or *radial* trapping

frequency $\omega_r = \omega_x = \omega_y$ ¹. For definiteness, we begin with the PGPE expressed in terms in dimensioned quantities (measured in, e.g., SI units). Denoting such quantities with a bar (e.g. \bar{x}), the PGPE is written

$$i\hbar \frac{\partial \bar{\psi}(\bar{\mathbf{x}})}{\partial \bar{t}} = \bar{\mathcal{P}} \left\{ \left(\bar{H}_{\text{sp}} + \bar{U}_0 |\bar{\psi}(\bar{\mathbf{x}})|^2 \right) \bar{\psi}(\bar{\mathbf{x}}) \right\}, \quad (4.1)$$

where \bar{H}_{sp} is the single-particle Hamiltonian

$$\bar{H}_{\text{sp}} = \frac{-\hbar^2 \bar{\nabla}^2}{2m} + \frac{m}{2} (\omega_x^2 \bar{x}^2 + \omega_y^2 \bar{y}^2 + \omega_z^2 \bar{z}^2) + i\hbar \bar{\Omega} \left(\bar{x} \frac{\partial}{\partial \bar{y}} - \bar{y} \frac{\partial}{\partial \bar{x}} \right), \quad (4.2)$$

and the projector

$$\bar{\mathcal{P}} \bar{f}(\bar{\mathbf{x}}) = \sum_{n \in \mathbf{C}} \bar{\phi}_n(\bar{\mathbf{x}}) \int d\bar{\mathbf{x}}' \bar{\phi}_n^*(\bar{\mathbf{x}}') \bar{f}(\bar{\mathbf{x}}'), \quad (4.3)$$

where the $\bar{\phi}_n(\bar{\mathbf{x}})$ are eigenstates of the dimensioned base Hamiltonian, i.e. $\bar{H}_0 \bar{\phi}_n(\bar{\mathbf{x}}) = \bar{\epsilon}_n \bar{\phi}_n(\bar{\mathbf{x}})$. We now introduce units

$$r_0 = \sqrt{\frac{\hbar}{m\omega_r}} \quad (4.4)$$

$$t_0 = \frac{1}{\omega_r} \quad (4.5)$$

$$\varepsilon_0 = \hbar\omega_r, \quad (4.6)$$

such that the dimensioned quantities are related to the dimensionless ones by $\bar{t} = t_0 t$, $\bar{\mathbf{x}} = r_0 \mathbf{x}$, etc. In particular, the general correspondence for wave functions is $\bar{\chi}(\bar{\mathbf{x}}) = r_0^{-3/2} \chi(\bar{\mathbf{x}}/r_0) \equiv r_0^{-3/2} \chi(\mathbf{x})$. We thus obtain the dimensionless PGPE

$$i \frac{\partial \psi}{\partial t} = \mathcal{P} \left\{ \left(H_{\text{sp}} + U_0 |\psi(\mathbf{x})|^2 \right) \psi(\mathbf{x}) \right\}, \quad (4.7)$$

where the dimensionless single-particle Hamiltonian has the form

$$H_{\text{sp}} = \frac{-\nabla^2}{2} + \frac{1}{2} (x^2 + y^2 + \lambda_z^2 z^2) + i\Omega \left(x \frac{\partial}{\partial y} - y \frac{\partial}{\partial x} \right), \quad (4.8)$$

with $\Omega \equiv \bar{\Omega}/\omega_r$ and $\lambda_z \equiv \omega_z/\omega_r$; the dimensionless interaction strength is

$$U_0 = \frac{\bar{U}_0}{\hbar\omega_r r_0^3}; \quad (4.9)$$

¹In chapters 7 and 9 we include an elliptical deformation of the trap in the xy plane. In this case ω_r refers to the ‘average’ radial trapping frequency, i.e. $\omega_r = \sqrt{(\omega_x^2 + \omega_y^2)/2}$ [130].

and the dimensionless projector \mathcal{P} is such that

$$\mathcal{P}f(\mathbf{x}) = \sum_{n \in \mathbf{C}} \phi_n(\mathbf{x}) \int d\mathbf{x}' \phi_n^*(\mathbf{x}') f(\mathbf{x}'), \quad (4.10)$$

where the dimensionless basis modes $\phi_n(\mathbf{x})$ are eigenstates of the dimensionless base Hamiltonian ($H_0\phi_n(\mathbf{x}) = \epsilon_n\phi_n(\mathbf{x})$) with dimensionless energy eigenvalues $\epsilon_n = \bar{\epsilon}_n/\varepsilon_0$.

Unit-normalised field

In most of this thesis we will work with a classical field $\psi(\mathbf{x})$ normalised to the total coherent-region population, $\int d\mathbf{x} |\psi(\mathbf{x})|^2 = \int d\bar{\mathbf{x}} |\bar{\psi}(\bar{\mathbf{x}})|^2 \equiv N_c$. However, in chapter 5 we consider the generic behaviour of the PGPE without explicit reference to physical units. In this chapter (only) we will therefore work in terms of a classical field which is normalised to unity, i.e., we define

$$\psi(\mathbf{x}) = \frac{\bar{\psi}(\bar{\mathbf{x}}) r_0^{3/2}}{\sqrt{N_c}}. \quad (4.11)$$

In that case the factor N_c is absorbed into the nonlinear coefficient, i.e., we have the PGPE

$$i \frac{\partial \psi(\mathbf{x})}{\partial t} = \mathcal{P} \left\{ (H_{\text{sp}} + C_{\text{nl}} |\psi(\mathbf{x})|^2) \psi(\mathbf{x}) \right\}, \quad (4.12)$$

where $C_{\text{nl}} = N_c \bar{U}_0 / \hbar \omega_r r_0^3$.

Quasi-two-dimensional system

In chapters 6-9 we will consider quasi-two-dimensional systems, which are physically three-dimensional systems tightly confined in the axial direction (see, e.g. [247, 55]). These are realised formally in our classical-field model by choosing a cutoff energy less than the oscillator spacing in the tightly confined (z) direction, i.e., $E_R < \hbar \omega_z$. We thus have

$$\mathcal{P}f(\mathbf{x}) = \sum_{n \in \mathbf{C}^\perp} \phi_n^\perp(x, y) \zeta_0(z) \int d^3\mathbf{x}' [\phi_n^\perp(x', y')]^* \zeta_0^*(z') f(\mathbf{x}'), \quad (4.13)$$

where the ground oscillator state $\zeta_0(z) = (\lambda_z/\pi)^{1/4} \exp(-\lambda_z^2 z^2/2)$ obeys

$$H_z \zeta_0(z) = \left[-\frac{1}{2} \frac{\partial^2}{\partial z^2} + \frac{1}{2} \lambda_z^2 z^2 \right] \zeta_0(z) = \frac{\lambda_z}{2} \zeta_0(z), \quad (4.14)$$

the set of functions $\{\phi_n^\perp(x, y)\}$ are the eigenstates of the transverse single-particle Hamiltonian $H_\perp = H_{\text{sp}} - H_z$ with eigenvalues $\{\epsilon_n^\perp\}$, and the region $\mathbf{C}^\perp \equiv \{n : \epsilon_n^\perp + \lambda_z/2 \leq E_R\}$.

Defining $\psi(\mathbf{x}) = \psi_\perp(x, y)\zeta_0(z)$, we therefore have

$$i\frac{\partial\psi_\perp(x, y)}{\partial t} = \mathcal{P}_\perp\left\{\left(H_\perp + \frac{\lambda_z}{2} + U_{2D}|\psi_\perp(x, y)|^2\right)\psi_\perp(x, y)\right\}, \quad (4.15)$$

where the effective two-dimensional interaction strength is

$$U_{2D} = \frac{\bar{U}_0}{r_0^2\hbar\omega_r\sqrt{2\pi}z_0} = \frac{\bar{U}_{2D}}{r_0^2\hbar\omega_r}, \quad (4.16)$$

where $z_0 = \sqrt{\hbar/m\omega_z}$ is the oscillator length in the tightly-confined dimension, and $\bar{U}_{2D} = \bar{U}_0/z_0\sqrt{2\pi}$ is the SI interaction strength appropriate to bosons in a quasi-2D geometry [205], and the two-dimensional projector is defined

$$\mathcal{P}_\perp f(x, y) \equiv \sum_{n \in \mathbf{C}^\perp} \phi_n^\perp(x, y) \int dx' dy' [\phi_n^\perp(x', y')]^* f(x', y'). \quad (4.17)$$

4.1.2 Spectral-Galerkin formulation

Now, as we have noted, the single-particle Hamiltonian can be separated as $H_{\text{sp}} = H_0 + \delta V(\mathbf{x}, t)$. As the projector \mathcal{P} is defined in terms of eigenfunctions of H_0 , we have immediately

$$[\mathcal{P}, H_0] = 0, \quad (4.18)$$

and so the equation of motion can be written as

$$i\frac{\partial\psi(\mathbf{x})}{\partial t} = H_0\psi(\mathbf{x}) + \mathcal{P}\left\{(\delta V(\mathbf{x}) + \lambda|\psi(\mathbf{x})|^2)\psi(\mathbf{x})\right\}. \quad (4.19)$$

Moreover, as the Hamiltonian H_0 is *diagonal* in the basis of eigenmodes $\{\phi_n(\mathbf{x})\}$, it is natural to take a *spectral Galerkin* approach, i.e., to derive from the PGPE (4.19) a set of coupled equations for the coefficients $\{\alpha_n\}$ in the expansion of ψ . This is achieved by calculating the overlap of each side with the general basis mode $\phi_m(\mathbf{x})$, i.e., by applying the operation $\int d\mathbf{x} \phi_m^*(\mathbf{x})$ to each side. As the projector \mathcal{P} is an identity operator for all basis modes ϕ_n , we have

$$\int d\mathbf{x} \phi_m^*(\mathbf{x}) \mathcal{P}\{f(\mathbf{x})\} = \int d\mathbf{x} \phi_m^*(\mathbf{x}) f(\mathbf{x}), \quad (4.20)$$

and thus

$$i\frac{\partial\alpha_m}{\partial t} = \epsilon_m\alpha_m + F_m[\psi(\mathbf{x})] + H_m[\psi(\mathbf{x})], \quad (4.21)$$

where

$$F_m[\chi(\mathbf{x})] = U_0 \int d\mathbf{x} \phi_m^*(\mathbf{x}) |\chi(\mathbf{x})|^2 \chi(\mathbf{x}), \quad (4.22)$$

$$H_m[\chi(\mathbf{x})] = \int d\mathbf{x} \phi_m^*(\mathbf{x}) \delta V(\mathbf{x}, t) \chi(\mathbf{x}). \quad (4.23)$$

We have used the functional notation to emphasise that the nonlinear term $F_m[\psi(\mathbf{x})]$ depends on the basis coefficients $\{\alpha_m\}$ through the whole projected field $\psi(\mathbf{x})$. By contrast, the perturbing potential term can be expressed as a matrix multiplication $H_m[\psi(\mathbf{x})] = \sum_n H_m[\alpha_n \phi_n(\mathbf{x})] = \sum_n H_{mn} \alpha_n$. However, in practice we will evaluate both terms using Gaussian quadratures, so we will not make this distinction, and we will in fact often group the two projected terms together as $G_m[\psi(\mathbf{x})] = F_m[\psi(\mathbf{x})] + H_m[\psi(\mathbf{x})]$.

4.1.3 Interaction-picture methods

The PGPE as written in equation (4.21) immediately lends itself to an ‘interaction picture’ approach [17]. The basic idea is to exploit the diagonal representation of the base Hamiltonian H_0 to transform the field (or more precisely its amplitude representation $\{\alpha_m\}$) to an interaction picture with respect to this operator. The effect of the remaining operators on the field is then evaluated efficiently in this interaction picture by traditional techniques for numerical integration, such as the *Runge-Kutta* family of integrators [215]. By eliminating the time-dependence of the field due to the base Hamiltonian H_0 from the problem, we are able to take much longer integration steps for the same level of numerical accuracy.

In this approach, the field is transformed to the interaction picture with respect to H_0 and relative to some time origin t_0 , so that

$$\psi^I(\mathbf{x}, t) = e^{i(t-t_0)H_0} \psi(\mathbf{x}, t) \equiv U_{\text{IP}} \psi(\mathbf{x}, t), \quad (4.24)$$

which yields the equation of motion

$$i \frac{\partial \psi^I(\mathbf{x}, t)}{\partial t} = \mathcal{P} \left\{ U_{\text{IP}} \left[\delta V + U_0 |U_{\text{IP}}^\dagger \psi^I(\mathbf{x}, t)|^2 \right] U_{\text{IP}}^\dagger \psi^I(\mathbf{x}, t) \right\}. \quad (4.25)$$

The spectral representation is again obtained by applying the operation $\int d\mathbf{x} \phi_m^*(\mathbf{x})$ to both sides of the equation, from which we obtain

$$i \frac{\partial \alpha_m^I}{\partial t} = G_m^I [\psi^I(\mathbf{x}, t)], \quad (4.26)$$

where

$$G_m^I[\chi(\mathbf{x}, t)] = \int d\mathbf{x} \phi_m^*(\mathbf{x}) U_{\text{IP}} \left[\delta V(\mathbf{x}, t) + U_0 |U_{\text{IP}}^\dagger \chi(\mathbf{x}, t)|^2 \right] U_{\text{IP}}^\dagger \chi(\mathbf{x}, t) \quad (4.27)$$

$$= e^{i\epsilon_m(t-t_0)} G_m \left[U_{\text{IP}}^\dagger \chi(\mathbf{x}, t) \right]. \quad (4.28)$$

This forms the basis for integration by a Runge-Kutta approach in the interaction picture, where the origin t_0 is judiciously chosen to minimise the number of interaction-picture transformations. In practice we use an *adaptive* Runge-Kutta variant [63], to control the accuracy of our solutions. This approach is briefly reviewed in appendix B.

4.1.4 Gaussian-quadrature evaluation of projected operators

In general, the basis modes we deal with are of the form

$$\phi_n(\mathbf{x}) = \mathcal{W}(\mathbf{x}) P_n(\mathbf{x}), \quad (4.29)$$

where $P_n(\mathbf{x})$ is a polynomial of some degree \mathcal{D}_{P_n} , and $\mathcal{W}(\mathbf{x})$ is some general *weight* function, which is common to all the basis modes $\phi_n(\mathbf{x})$. For example, in the case of the Hermite basis modes of the one-dimensional harmonic oscillator, $\mathcal{W}(x) = e^{-x^2/2}$. Consequently, the classical field can be expressed

$$\psi(\mathbf{x}) = \sum_{n \in \mathbf{C}} \alpha_n \phi_n(\mathbf{x}) = \mathcal{W}(\mathbf{x}) Q(\mathbf{x}), \quad (4.30)$$

where $Q(\mathbf{x})$ is a polynomial of *at most* some degree \mathcal{D}_Q , which is set by the cutoff energy and thus the allowed basis indices $n \in \mathbf{C}$. We now restrict our attention to the case $\delta V(\mathbf{x}, t) = 0$, and consider the evaluation of the nonlinear matrix element $F_m[\psi(\mathbf{x})]$

$$F_m[\psi(\mathbf{x})] = U_0 \int d\mathbf{x} \phi_m^*(\mathbf{x}) \mathcal{W}^3(\mathbf{x}) |Q(\mathbf{x})|^2 Q(\mathbf{x}) \quad (4.31)$$

$$= U_0 \int d\mathbf{x} \mathcal{W}^4(\mathbf{x}) R_m(\mathbf{x}), \quad (4.32)$$

where $R_m(\mathbf{x})$ is a polynomial of degree at most \mathcal{D}_{R_m} . The relevant case is that in which $\mathcal{W}^4(\mathbf{x}) \equiv W(\mathbf{x})$, where $W(\mathbf{x})$ is the *weight factor* of a Gaussian quadrature rule [215]

$$\int d\mathbf{x} W(\mathbf{x}) f(\mathbf{x}) \approx \sum_{j=1}^{N_Q} w_j f(\mathbf{x}_j), \quad (4.33)$$




Figure 4.1: Visual representation of the Gaussian-quadrature method for evaluating the projected nonlinear term. (Figure adapted from reference [28].)

for approximating the integral. The quadrature *abscissas* \mathbf{x}_j form a spatial grid which is in general not uniformly spaced, and the quadrature *weights* w_j are the corresponding numeric weighting values appropriate to the quadrature rule. The crucial point is that given the bound \mathcal{D}_{R_m} on the degree of polynomial $R_m(\mathbf{x})$, the order N_Q of the quadrature rule (4.33) can be chosen such that the integral $F_m[\psi(\mathbf{x})]$ can be evaluated *exactly*:

$$F_m[\psi(\mathbf{x})] = U_0 \sum_{j=1}^{N_Q} w_j R_m(\mathbf{x}_j). \quad (4.34)$$

To use this approach in practice, the roots \mathbf{x}_j and weights w_j appropriate to the quadrature rule of the requisite order are precomputed before the evolution begins, as is the *transformation matrix* $U_{kn} \equiv \phi_n(\mathbf{x}_k)$ between the spectral and quadrature-grid representations of the classical field. The general procedure for evaluating $F_m[\psi(\mathbf{x})]$ is then (figure 4.1):

1. Transform from the coefficient representation to a spatial representation on the quadrature grid

$$\psi(\mathbf{x}_k) = \sum_{n \in \mathbf{C}} U_{kn} \alpha_n. \quad (4.35)$$

2. Construct the integrand on the quadrature grid by local multiplication

$$R(\mathbf{x}_k) = w_k |\psi(\mathbf{x}_k)|^2 \psi(\mathbf{x}_k). \quad (4.36)$$

3. Obtain the desired matrix elements by performing the inverse transformation

$$F_m[\psi(\mathbf{x})] = U_0 \sum_{k=1}^{N_Q} U_{km}^* R(\mathbf{x}_k). \quad (4.37)$$

4.2 Performance of the Gauss-Laguerre algorithm

In this section we consider the numerical performance of the Gauss-Laguerre-quadrature algorithm for the integration of the rotating-frame PGPE. The efficiency of the algorithm is a crucial concern, and indeed limits the size of system and duration of evolution we can simulate in a given period of real (wall) time. The execution time of numerical algorithms depends in general both on the speed of the processor, and the bandwidth of the memory in the computer on which it is executed. All the results we present in this section are for the C-optimised MATLAB implementation of the algorithm², timed on workstation with a 2.66GHz Intel Xeon processor and 667MHz DDR2 RAM, which is of comparable performance to the compute nodes of the cluster on which the majority of production simulations were run³.

4.2.1 Base rotating-frame algorithm

We begin by considering the performance of the base algorithm developed by Bradley [34, 37]. In this algorithm (section C.3.1), the quadrature grid (cf. figure 4.1(b)) is a set of polar coordinates $\{r_k, \theta_j\}$, with N_x points in the radial direction and N_θ points in the azimuthal direction. In general, both the multiplicity \mathcal{M} of the basis modes (cf. figure 4.1(a)) and the number of points in each direction on the quadrature grid will depend both on the cutoff energy and the angular velocity of the rotating frame in which the projection is effected. As noted by Bradley, by considering the operations involved in the evaluation of the nonlinear term (cf. [71]), one finds that the computational load of the evaluation of the nonlinear term by Gauss-Laguerre quadrature should scale as $O(NN_x \ln N_\theta)$, where $N = \max(N_x, \ln N_\theta)$ [37]⁴. As the evaluation of this term forms the most numerically intensive part of the Runge-Kutta algorithms, we expect similar scaling for the total integrator. Here, we consider the ‘real world’ performance of the algorithm,

²In order to increase the efficiency of the algorithm, we have written its core processor-intensive operations (i.e., matrix multiplications and calls to the fast Fourier transform library [96]) in the C language, using MATLAB’s MEX interface [179].

³The actual performance on the multi-core compute nodes will depend on the system load conditions during execution, and so we consider the performance on an isolated workstation for clarity.

⁴In fact for nontrivial cutoff parameters we always have $N_x > \ln N_\theta$, and so the algorithmic scaling becomes simply $O(N_x^2 N_\theta)$.

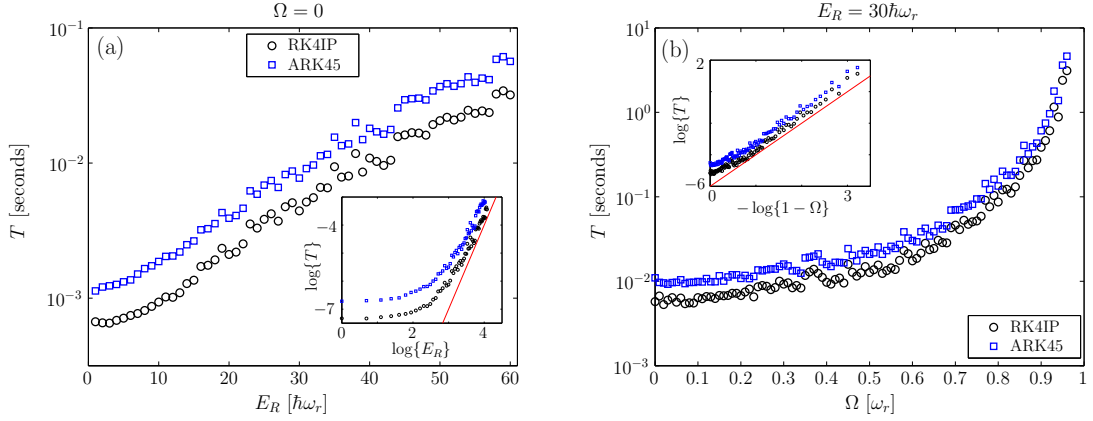


Figure 4.2: Measured step times for RK4IP and ARK45-IP implementations of the rotating-frame PGPE. a) Dependence of the step time T on the cutoff energy E_R at fixed rotation rate $\Omega = 0$. The inset log-log plot shows that the algorithm approaches the expected $O(E_R^3)$ scaling as E_R increases. b) Dependence of the step time on the angular velocity Ω of the rotating frame at fixed cutoff energy $E_R = 30\hbar\omega_r$. The inset log-log plot shows that the algorithm approaches a scaling $O((1 - \Omega)^{-2})$ as $\Omega \rightarrow 1$.

including the dependence of the integrator step time on the energy cutoff and frame-rotation frequency, and the difference between the fixed-step-size RK4IP implementation and the adaptive ARK45-IP variant (appendix B).

Dependence on energy cutoff

We first consider the average step time of the two Runge-Kutta algorithms as a function of the cutoff energy E_R , for the nonrotating case $\Omega = 0$. For each value of the cutoff energy we plot in figure 4.2(a) the mean of 100 measured step times for both the RK4 (circles) and ARK (squares) algorithms. In both cases the scaling of the step time T with cutoff energy E_R is comparatively flat for low cutoffs $E_R \lesssim 10\hbar\omega_r$, as the step time is dominated by fixed overhead factors in this small-basis limit. For larger quadrature grids, the step-time scaling is reasonably close to the $O(E_R^3)$ scaling we expect⁵, as indicated in the inset to figure 4.2(a), in which we plot $\log(T)$ against $\log(E_R)$. We include a line $y = mx + c$ with slope $m = 3$ for comparison. We note that the ARK steps take on average 1.8 times as long to execute as the RK4 steps, which is reasonable given that six evaluations of the nonlinear term are required for each step of the ARK algorithm, versus four evaluations for the RK4 algorithm, but also suggests that a significant amount of time is involved in the additional diagonal multiplications and copying operations in the ARK implementation.

⁵In the nonrotating limit $\Omega = 0$, we have $N_x = \bar{N} + 1$ and $N_\theta = 4\bar{N} + 1$, thus $NN_xN_\theta = N_x^2N_\theta \approx 4\bar{N}^3$.

We also note that the increase in step time with the cutoff height is not monotonic. Indeed it is clear from the fact that ‘glitches’ in the upwards progression of step times appear simultaneously in both the RK4 and ARK data sets that these variations are reproducible, and not merely statistical fluctuations. This non-monotonicity is a consequence of the fast Fourier transform algorithms used to evaluate the azimuthal quadrature⁶, for which the efficiency of the transform depends not only on the number N_θ of azimuthal grid points, but on the factorisation of this quantity as a product of prime numbers. In particular, quadrature grids for which the number N_θ of azimuthal quadrature points is prime can be significantly slower than comparable grids with more favourable factorisations.

Dependence on rotation rate

As the rotation frequency of the projector is increased, the coherent region of the classical field becomes increasingly skewed towards modes with higher angular momentum projections. The size of quadrature grid required to evaluate the nonlinear integral increases dramatically (in both the azimuthal and radial dimensions) as more high angular momentum modes are included in the calculation. We consider here the dependence of the step time on the rotation frequency Ω of the frame in which the PGPE solution is propagated, by comparing the execution times of step evaluations at fixed energy cutoff $E_R = 30\hbar\omega_r$ (figure 4.2(b)). In general we expect the step time to scale like⁷ $T \propto (1 - \Omega)^{-3}$. However in practice we find $T \propto (1 - \Omega)^{-2}$, as shown in the inset, where we plot $\log(T)$ against $-\log(1 - \Omega)$, and include a line with slope $m = 2$ for comparison. It is not clear why we observe this scaling (which is *more favourable* than the theoretical prediction), rather than the expected $T \propto (1 - \Omega)^{-3}$. It is possibly due to some compiler optimisation arising in the multiplication of the triangular transformation matrices involved in the quadrature rule, which have many zero entries and become increasingly skewed as the rotation rate Ω increases.

4.2.2 Perturbing potential term

The significant extension of the Gauss-Laguerre quadrature algorithm employed in this thesis over that used in other work [37] is the addition of the elliptical potential term (appendix C). It is thus of interest to quantify how the inclusion of this term affects the step times. In this section we compare the ‘base’ algorithm (with ellipticity $\epsilon = 0$) against the algorithm with $\epsilon \neq 0$. For simplicity, we only compare results for the adaptive (ARK) algorithm, this being the algorithm we use for all simulations in practice. Plotting the

⁶In practice we use the FFTW library to evaluate discrete Fourier transforms [96, 97].

⁷We have $N_x \approx \bar{N}/(1 - \Omega)$ and $N_\theta = 2(\bar{l}_- + \bar{l}_+) + 1 \approx 4\bar{N}/(1 - \Omega)$, thus $N_x^2 N_\theta \approx 4\bar{N}/(1 - \Omega)^{-3}$.

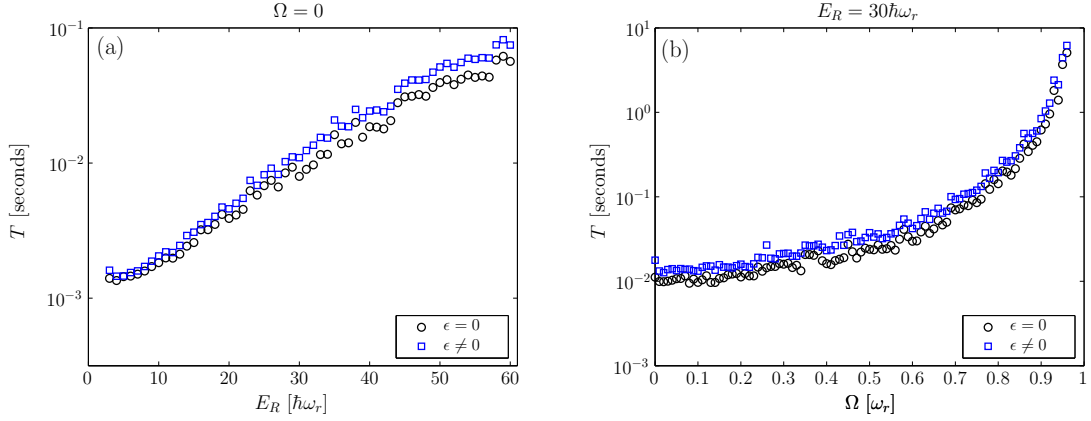


Figure 4.3: Effect of the perturbing potential term $\delta V(\mathbf{x})$. a) ARK45-IP step times with and without the additional quadrature rule for the perturbing trap anisotropy, as a function of cutoff energy E_R at fixed rotation rate $\Omega = 0$. b) Measured step times as a function of rotation rate Ω at fixed cutoff energy $E_R = 30\hbar\omega_r$.

step times in the presence of the perturbation alongside those obtained in its absence against the cutoff energy E_R with $\Omega = 0$, (figure 4.3(a)), we see that the algorithm is fractionally slower, with the step time increasing from $\approx 5\%$ to $\approx 40\%$ of that of the base algorithm, over the range of cutoff heights considered. Plotting the dependence of step times on the rotation rate Ω (figure 4.3(b)), we observe that the step time associated with the perturbing potential is $\approx 40\%$ of the base step time for all Ω we consider, in agreement with the $\Omega = 0$ values obtained for this cutoff height ($E_R = 30\hbar\omega_r$).

4.3 Comparison to other discretisation schemes

As we have already noted, in general any numerical implementation of GP-like classical field evolution must involve some discretisation of the continuum field. In this section we compare the projected approach to other common discretisation schemes applied to the integration of the GP equation. We discuss the various pathologies and ambiguities that arise in the common Cartesian-grid discretisation schemes, and argue that the projected-GPE approach we adopt in this thesis is the unique approach which is free of these issues. This is of particular importance as in chapters 8 and 9 we contrast the results of our PGPE calculations with those of other simulations which use Cartesian-grid methods.




Figure 4.4: Phase-space region \mathcal{D}_{ho} covered by the first $\bar{n} + 1 = 11$ harmonic oscillator modes (circles). The square indicates the region \mathcal{D}_{pw} enclosed by the corresponding optimal plane-wave basis with $\bar{n} + 1$ grid points (see [36]). Dashed circles indicate harmonic oscillator modes which are not accurately resolved in the plane-wave basis.

4.3.1 Phase-space covering

One-dimensional harmonic oscillator

In this section we follow Bradley *et al.* [36] in considering the phase space of a one-dimensional harmonic oscillator (of frequency ω), bounded by an energy cutoff $E_c = (\bar{n} + 1/2)\hbar\omega$. In terms of appropriately scaled quantities x and k (see section 4.1.1)⁸, the phase space is thus bounded by a circle, as illustrated in figure 4.4.

In the projected formalism, the restriction of the system to this region of phase space is effected *exactly* by restricting the basis of the field to the first $\bar{n} + 1$ harmonic oscillator modes. An alternative discretisation of the system is in terms of a regularly spaced grid in position space. In this case, there are now *two* cutoffs we must choose: the spatial extent L of the grid, and the spacing Δx between grid points, which determines the *Nyquist wave number* $k_c = \pi/\Delta x$, which is the magnitude of the largest wave vector resolvable on the grid. Alternatively, we can determine the grid completely by the number of points N_x and the grid extent L . We note that the x -space and k -space limits of the grid are independent, i.e., the same wave-number cutoff k_c applies for all spatial positions $|x| \leq L/2$, and so the corresponding region of phase space spanned by this plane-wave basis is a *rectangle* in

⁸The units of momentum and wave number corresponding to the system of units introduced in section 4.1.1 are $p_0 = \sqrt{\hbar m \omega_r}$ and $k_0 = 1/r_0 = p_0/\hbar$, respectively. In this system of units momentum and wave vector are thus equivalent, and so we use the two concepts interchangeably in this section.

phase space. It is important to note that, as the wave-number cutoff $k_c = \pi N_x/L$, the (2D) phase-space ‘volume’ spanned by the rectangular region $V_{\text{pw}} = 2\hbar k_c L = hN_x$; i.e., each mode corresponds to a region of volume h , as is the case for the harmonic oscillator basis, and as it should be for the quantum system (Weyl’s law [18]). Bradley *et al.* [36] showed that for a given number of points N_x , the grid with extent $L_{\text{opt}} = \sqrt{2\pi\hbar N_x/m\omega}$ is *optimal*, in the sense that the boundary of this phase space is *square* (in the appropriate units), and is thus the closest match to the circular phase-space region spanned by the corresponding harmonic basis. This region is indicated by its square boundary in figure 4.4. We note that the energy corresponding to the extremes of this region in each individual direction is $E_c^{(\text{pw})} = (m/2)\omega^2 L^2 = \hbar^2 k_c^2/2m = (\pi/4)N_x\hbar\omega$.

The authors of [36] then proceeded to diagonalise the single-particle Hamiltonian of the harmonic oscillator in both the harmonic basis, and the plane-wave basis with an equal number of modes ($N_x = \bar{n} + 1$). In the harmonic basis this is trivial, as the basis is, by definition, comprised of the first $\bar{n} + 1$ oscillator modes. By contrast, the diagonalisation of the Hamiltonian in the plane-wave basis accurately reproduces the lowest modes, but with increasing energy the modes become less well resolved in this basis. As the eigenmodes of the Hamiltonian in this basis essentially correspond to variational approximations to the true modes subject to the ‘ansatz’ that they are comprised of linear combinations of the N_x plane-wave modes, it is intuitive that the energies of these modes increasingly overestimate the true harmonic-oscillator mode energies. In the example of $N_x = \bar{n} + 1 = 16$ considered in [36], approximately the top one-third of modes were poorly represented in the *optimal* plane-wave basis, while other choices of plane-wave basis in general did a worse job of representing the modes.

In general, it is hard to quantify exactly what effect such error in the mode energies has on the classical-field dynamics. Some insight can be gained by using the principle of equipartition (section 3.3.2). At first sight, one might consider that a change in the energies would have little overall effect on the thermal equilibrium, as for a given total energy E , equipartition predicts that $E = \sum_j \epsilon_j n_j = \mathcal{M}k_B T$, so that the temperature is determined solely by the number of modes \mathcal{M} and the total thermal energy E , and does not depend on the actual values of the particular mode energies $\{\epsilon_j\}$. However for a fixed total energy (and thus temperature), the individual mode occupations n_j must vary in inverse proportion to the energies ϵ_j , and so the total thermal population $N = \sum_j n_j$ will vary with the $\{\epsilon_j\}$. We therefore conclude that while the relation between energy and temperature in general depends on the number of classical-field modes \mathcal{M} [170], the relationship between energy, temperature and *chemical potential* will in general depend on the precise energies of the modes, and will therefore be corrupted for non-ideal discretisations of the classical field.

The arguments here simply consider the single-particle modes, i.e., the thermal equilibrium of a noninteracting system. In general we are of course concerned with the thermal equilibrium of an interacting system. To the extent that a Bogoliubov description of the normal modes is valid, the relevant issue is the accuracy with which the normal modes are resolved within the propagation basis. Again, to the extent that the Bogoliubov modes return to a single-particle description at high energies, we expect the most accurate representation to be obtained in the appropriate harmonic-oscillator basis. However, even in this basis we must recognise that the highest-energy Bogoliubov modes will not be well represented in the harmonic basis, and in chapter 9 we show that a signature of this breakdown of the mode basis can be observed in the results of an interacting classical-field simulation. Nevertheless, as the field we evolve exists in a single-particle space, the basis of single-particle energy eigenstates is the optimal choice, and we must accept some error in the description of the highest normal modes of the interacting system as the minimal penalty for effecting the cutoff necessary for a classical-field treatment of the low-energy dynamics of the field.

Higher dimensionalities

We have shown that in one spatial dimension, where the corresponding classical phase space is two-dimensional, we can reasonably well approximate the boundary of phase space defined by a cutoff in single-particle (harmonic-oscillator) energy by a plane-wave (equidistant grid) discretisation. In higher dimensionalities this is no longer the case, as we now discuss. In one dimension, defining the cutoff by a single parameter (the energy) leads to the result that the momentum cutoff is *local*, i.e., the cutoff wave number is a function of the position coordinate x . For example, for a cutoff defined in terms of energy of a harmonic oscillator, the range of available momentum components below the cutoff is greatest at the centre of the trap, and smoothly vanishes as $|x|$ approaches the *classical turning point* associated with the cutoff energy. In higher dimensionalities, defining a cutoff in terms of the energy means that the cutoff in each of the $2 \times d$ dimensions of phase space is dependent on the coordinates in *all* the other dimensions. Thus, in two spatial dimensions, in a harmonic basis, we can define a cutoff in terms of an energy $E_c = (\bar{n} + 1)\hbar\omega$, leading to the constraint on the excitation numbers of the x and y oscillators $n_x + n_y \leq \bar{n}$, yielding a total of $\mathcal{M}_{\text{ho}}^{2\text{D}} \approx \bar{n}^2/2$ modes (corresponding to a phase-space volume $V_{\text{ho}}^{2\text{D}} \approx \bar{n}^2 h^2/2$). It is not clear how to best approximate this region of phase space with a plane-wave cutoff. One approach is to take $N_x = N_y = n$, with $L_x = L_y = L_{\text{opt}}(N_x)$, so that the boundary of the plane-wave space most closely matches that of the harmonic-oscillator space in the $k_x x$ and $k_y y$ planes. This yields a total number

of modes $M_{\text{pw}}^{2\text{D}} = \bar{n}^2$ (corresponding to phase-space volume $V_{\text{pw}}^{2\text{D}} = \bar{n}^2 h^2$). This plane-wave basis thus contains *twice* as many modes as the harmonic basis. A simple equipartition argument then predicts that the equilibrium temperature will be $\sim 1/2$ of the correct value, in the plane-wave case. The reader can easily deduce that in three dimensions, a similarly constructed plane-wave discretisation now encloses approximately *six* times the volume of the appropriate harmonic-basis discretisation, leading to an equilibrium temperature one-sixth of that of the ideal harmonic-basis description⁹. Assuming (as in the example of [36]) that the top one-third of harmonic oscillator modes in each direction are spuriously resolved under the plane-wave cutoff implies that the top $\approx 10\%$ of the 3D-oscillator modes are poorly resolved, with overestimated energies.

In summary, in three-dimensional calculations, the plane-wave basis which ‘best’ matches the (hyperspherical) energy-cutoff surface in phase space contains *six* times the correct number of modes, a significant fraction of which are seriously ill-resolved in the available phase space, and have spuriously high energies. Any larger ‘box’ in phase space would include yet more spurious degrees of freedom, while the use of any smaller box would lead to severe corruption of a greater fraction of the modes contained in it. It is therefore clear that the thermodynamics of the system would be strongly corrupted by this poor representation of the phase space, regardless of the precise dimensions of the plane-wave grid. An important aspect of a classical-field defined in terms of single-particle *energy eigenmodes* in two and three dimensional calculations is therefore that the resulting cutoff properly accounts for the local shape of the energy surface in all dimensions of phase space, as this surface can not be well approximated by the cutoff surface resulting from a plane-wave or Cartesian-grid basis.

4.3.2 Implementing spatial derivatives

Thus far we have simply discussed how the phase space defined by a Cartesian grid compares with the phase space spanned by the harmonic oscillator basis. Here we consider in more detail how the spatial integration of a GP-like partial differential equation is implemented, i.e., how the various operators of the GP equation are evaluated on the computational grid. In the case of a regular Cartesian grid in position space, the potential and nonlinear terms are *local*, and thus trivial to evaluate, and the difficulty arises in evaluating *derivatives* of the field, i.e., the Laplacian $-\nabla^2$ appearing in the kinetic-

⁹Note that reversing these arguments one finds immediately that the plane-wave basis chosen to most accurately match the *volume* of the harmonic-oscillator phase space will result in dramatically smaller cutoffs in each dimension of phase space, and will therefore unnaturally constrain the physics of the system.

energy term, as well as others that may result from the transformation to a translating or rotating frame.

Finite-difference and Fourier derivatives

One approach to calculating derivatives of the classical field on a position-space grid is by the use of *finite-difference* approximations to the relevant derivative operator. A thorough account of such finite-difference approximations is given by Blakie [26], and such approximations are employed in (e.g.) the Crank-Nicolson method for evolving (nonlinear) Schrödinger equations (see, e.g., [192]). Usually one approximates derivatives by *central* finite difference stencils. For example, indexing the discretised field in terms of the grid points at which it is evaluated, $\psi_i \equiv \psi(x_i)$, the simplest central finite-difference approximation to the derivative $\partial_x \psi$ is given by [26]

$$(\delta_x \psi)_i = \frac{1}{2\Delta x} [\psi_{i+1} - \psi_{i-1}]. \quad (4.38)$$

The key point is that $(\delta_x \psi)_i$ is a good approximation to $\partial_x \psi$ at x_i for variations of ψ on scales *large* compared to Δx , or equivalently, for wave numbers much smaller than the *cutoff* (Nyquist) wave number k_c . However, this leaves a large number of intermediate wave vectors for which the derivatives are *not* accurately evaluated. The most important example is provided by the approximation to the Laplacian $-\nabla^2$, which we here consider in this low-order scheme, in one dimension where it becomes [192]

$$(-\delta_x^2 \psi)_i = -\frac{1}{(\Delta x)^2} [\psi_{i+1} - 2\psi_i + \psi_{i-1}]. \quad (4.39)$$

If we assume periodic boundary conditions (so that $\psi_{N_x+1} = \psi_1$), then in the basis of position eigenstates $\zeta_j(x_i) = \delta_{ij}$ the operator $-\delta_x^2$ is a sparse (band-diagonal) circulant matrix which is symmetric and has a complete basis of eigenvectors, which are the plane waves $\phi_i = e^{\pm i k_i x}$ (or equivalently the pairs $\sin(k_i x)$ and $\cos(k_i x)$), where the k_i are the momentum modes resolvable on the Cartesian grid. These are of course the correct eigenvectors, in the sense that they are eigenstates of the continuum Laplacian $-\nabla^2$ we wish to approximate. However, their corresponding eigenvalues differ from those of $-\nabla^2$. In fact, we find that $-\delta_x^2 e^{i k_i x} = \lambda_i e^{i k_i x}$ where the eigenvalues

$$\lambda_{k_i} = \frac{2}{(\Delta x)^2} \left[1 - \cos \left(\frac{\pi k_i}{k_c} \right) \right]. \quad (4.40)$$

Clearly for wave numbers which are *small* compared to k_c , this dispersion relation reduces to the continuum result $\lambda_k = k^2$. However, for k as small as $k_c/2$ the error in the eigenvalue

λ_k is of order 20%. In fact, the dispersion relation (4.40) exhibits the peculiar feature of *negative mass* for excitations with $k > k_c/2$, as is well known for lattice systems (e.g., in solid-state physics)¹⁰.

An alternative definition of the Laplacian on the Cartesian grid is obtained from the fundamental relation [63]

$$-\nabla^2\psi(x) = \int dk k^2 e^{-ikx} \times \frac{1}{2\pi} \int dy e^{iky} \psi(y). \quad (4.41)$$

The analogous discrete form of this relation (written in terms of discrete Fourier transforms) has *exact* eigenvalues $\lambda_{k_i} = k_i^2$. Thus while both approximations of the Laplacian approach the same continuum limit as $\Delta_x \rightarrow 0$ and thus $k_c \rightarrow \infty$, only the Fourier-defined Laplacian yields the ‘correct’ eigenvalues throughout the computational domain. Of course one expects the accuracy of the finite-difference Laplacian to be improved by using higher-order finite-differencing stencils, and we ultimately expect convergence to the Fourier-defined Laplacian in the limit that the order of the stencil approaches the total number of points on the grid. In this limit, however, the finite-difference Laplacian is completely nonlocal, and acts as a *dense* matrix on the configuration space. Although the Fourier-transform operator is fundamentally also a dense matrix on the configuration space, *fast Fourier transform* algorithms allow for the application of the transform using $O(N \log N)$ floating-point operations, as compared to $O(N^2)$ for a standard matrix multiplication, so the numerical advantage of using the Fourier-defined Laplacian is obvious.

Aliasing

Another important concern with the use of grid methods is that of momentum-mode aliasing. This is so named because momentum modes e^{ikx_i} can only be properly resolved on the grid if the magnitude of the wave vector is not greater than the Nyquist wave number ($|k| \leq k_c$). Mathematically, the plane wave with $k = k_j + 2nk_c$ is indistinguishable from that with $k = k_j$ on the grid $\{x_i\}$, as both waves have the same values on all points on the grid. Thus the momentum mode k_j acts as an *alias* for all modes $k_j + 2nk_c; n \in \mathbb{Z}$. In digital signal processing, one must therefore be aware that any frequency components above the Nyquist cutoff will be indistinguishable from their aliases below the cutoff. In the study of nonlinear Schrödinger equations, the issue is that the nonlinear potential $U_0|\psi|^2$ mixes momentum components, and thus (potentially) attempts to generate mo-

¹⁰Note that the issue here is that of a poor definition of the operator, not a restriction on the space on which it acts (as we considered in the previous section), so the eigenvalues need not be overestimated with respect to their ‘true’ values.

momentum components which are not resolvable on the grid, resulting in spurious coupling to their mode aliases.

A common misconception is that this phenomenon is introduced by the explicit use of the Fourier transform in GP propagation algorithms. This is not the case, and the spurious momentum-nonconserving processes simply result from the action of the (local) nonlinear potential attempting to populate unresolvable modes. To demonstrate this we show that the states are spuriously connected by the nonlinear potential. Usually a potential V causes a transition from a state $|\zeta_a\rangle$ to state $|\zeta_b\rangle$ if the ket $V|\zeta_a\rangle$ has some component along $|\zeta_b\rangle$. The behaviour resulting from the nonlinear potential $|\psi|^2$ can similarly be deduced, except that the potential of course now depends on the ket it is applied to. We consider the wave function consisting of two momentum modes k_a and k_b , defined on the $\{x_i\}$ grid, $\psi(x_i) = ae^{ik_ax_i} + be^{ik_bx_i}$. Acting on this state with the *local* potential $V_{\text{nl}}(x_i) = |\psi(x_i)|^2$, yields

$$\begin{aligned} [V_{\text{nl}}\psi](x_i) &= a^2b^*e^{i(2k_a-k_b)x} + a^*b^2e^{i(2k_b-k_a)x} \\ &\quad + (|a|^2 + 2|b|^2)ae^{ik_ax_i} + (2|a|^2 + |b|^2)be^{ik_bx_i}, \end{aligned} \quad (4.42)$$

thus the nonlinear potential transfers population to the plane wave $e^{i(2k_a-k_b)x}$. In the particular ‘worst case’ considered by Norrie [189]: $k_a = -k_c$, $k_b = k_c - \Delta k$, we see that the nonlinear potential attempts to transfer population to the mode $-3k_c + \Delta k$, resulting in spurious transfer to the alias mode $-k_c + \Delta k$.

We note that grid methods which are *not* based on Fourier-transform evaluation of derivatives are frequently implemented with *Dirichlet* boundary conditions, i.e., the wave function is constrained to specified values (e.g., zero) at the extremes of the grid [192]. One can easily show that a wave function consistent with such a constraint (i.e., expanded in terms of $\cos(k_ix)$ and $\sin(k_ix)$) suffers from the same aliasing problem under the action of the local nonlinear potential.

We note that the Fourier-based grid methods are thus more self-consistent than finite-difference-based methods in that they explicitly take into account the periodicity in both x and k space, the k -space periodicity being an implicit feature of all grid methods, due to the mode aliasing. Ultimately, the problem with Fourier-based methods is that the phase space spanned by the plane-wave basis modes is not only geometrically ill-formed, but is also topologically incorrect: due to the periodic nature of the position and momentum coordinates, the phase space has no real boundary, and is instead a toroidal structure. Momentum-mode aliasing can be eliminated for Fourier-based simulations of homogeneous systems with periodic boundary conditions in space [63] and simulations in free space [189], by a projected generalisation of the method. However, there is no simple

way to accurately match the correct phase space of (e.g.) a harmonically trapped system using a plane-wave basis.

Spectral and discrete-variable-representation methods

An alternative approach to the discretisation of the continuous nonlinear Schrödinger equation is through the use of a *spectral representation*. In this approach, one expands the Schrödinger wave function over a set of basis functions with known derivatives, so that the action of (e.g.) the Laplacian can be evaluated exactly. The difficulty then is in evaluating the matrix elements of potential terms between these basis functions. In a *discrete variable representation* [228], the basis functions are chosen such that an associated quadrature rule exists, and the functions form a complete basis for functions evaluated at the quadrature points $\{x_k\}$. A transformation between the basis expansion and a set of corresponding *position eigenfunctions* then allows for accurate evaluation of operators which are diagonal in position space, while retaining a simple matrix representation for the Laplacian.

The authors of [71] presented a technique for evolving the time-dependent GPE for a harmonically confined condensate in a basis of the Hermite-Gaussian eigenmodes of the harmonic-oscillator Hamiltonian. Such a *spectral-Galerkin* approach combines the accuracy of the Gaussian quadrature rule for the evaluation of the nonlinear potential, while *exactly* diagonalising the whole single-particle Hamiltonian. The Gauss-quadrature methods presented in appendix C are a refinement of this approach, with the distinction that the precise set of basis modes used is carefully chosen to *define* the classical-field system, according to a prescribed single-particle energy cutoff, and to ensure the optimal representation of all modes with energies below this cutoff.

Chapter 5

Anomalous moments in classical-field theory

Equilibrium classical-field methods make no assumptions about the fluctuation statistics of the atomic field beyond the neglect of *quantum* fluctuations. In high-temperature regimes, they therefore *implicitly* describe all relevant physical processes that shape the correlations of the field. This is in stark contrast to self-consistent mean-field theories (section 2.2.3), in which one *begins* with assumptions about the fluctuation statistics of the Bose field: the separation of the field into condensed and noncondensed (thermal) components, and the factorisation of the moments of the thermal component. Although classical-field methods provide an accurate, nonperturbative description of the field, due to the lack of such constraining assumptions we frequently face difficulties in *interpreting* the results that they describe, and specifically in relating them to more traditional descriptions of Bose-Einstein condensation, such as the mean-field theories.

In this chapter, we demonstrate that meaningful information is encoded in the *temporal* behaviour of classical-field trajectories. We show that equilibrium classical-field trajectories exhibit a kind of symmetry-breaking phenomenon, in the sense that the condensate mode exhibits a well-defined phase at any point in time. Due to the phase symmetry of the classical-field Hamiltonian, the condensate phase diffuses over time [235], restoring the symmetry. However, averages of the classical field in an appropriate phase-rotating frame taken on time scales *short* compared to that of the phase diffusion yield finite values for moments of the field which are formally zero in its microcanonical density. These moments correspond to the *anomalous moments* of the Bose field in symmetry-breaking mean-field theories. Applying this methodology to PGPE simulations of the harmonically trapped Bose gas, we show that the first moment of the field obtained in this approach agrees well with the condensate defined by the Penrose-Onsager criterion based on one-body correla-

tions (section 3.3.1), over a wide range of temperatures. We demonstrate the generality of this prescription for evaluating anomalous moments of the classical-field variables by calculating the anomalous thermal density of the field, which we find to have form and temperature dependence consistent with the results of mean-field theory calculations.

This chapter is organised as follows: in section 5.1 we introduce the PGPE system we study in this chapter. In section 5.2 we discuss how a nonzero first moment of the classical field arises from its temporal coherence, and make a quantitative comparison to the condensate defined by the Penrose-Onsager measure of one-body (spatial) coherence. In section 5.3 we consider the dependence of this first moment and its associated frequency on the energy of the classical field. In section 5.4, we consider the anomalous second moments which comprise the classical-field *pair matrix*, and construct the anomalous thermal density of the field. In section 5.5 we summarise and present our conclusions.

5.1 System

In this chapter we consider the behaviour of a nonrotating, three-dimensional harmonically trapped classical field. Our interest here is in characterising the correlations of the field, rather than to provide a direct comparison to an experimental system, and we therefore study a *unit-normalised* classical field, which evolves according to equation (4.12). We take the trap anisotropy parameter $\lambda_z = \sqrt{8}$, corresponding to a typical, moderately oblate three-dimensional trap geometry. We choose the nonlinearity parameter $C_{\text{nl}} = \sqrt{2} \times 500$. The corresponding ground (Gross-Pitaevskii) eigenstate of this system thus has energy $E_g \approx 9N_c\hbar\omega_r$, and GP eigenvalue $\lambda_g \approx 12\hbar\omega_r$. Using the criteria discussed in section 3.2 we therefore set the energy cutoff $E_R = 31\hbar\omega_r$. We form initial states with energies in the range $E \in [9.5, 24.0]N_c\hbar\omega_r$, by in each case mixing an admixture of a random excited state with the Thomas-Fermi ground state of the unit-normalised field [29, 28]. We then generate equilibrium states with a range of temperatures by evolving these randomised initial states in real time, exploiting the ergodicity of the PGPE (see section 3.3). In practice we find that the classical-field configurations are well equilibrated after $120\omega_r^{-1}$ of real-time evolution, and we perform our analysis on the subsequent evolution of the field.

5.2 Temporal coherence and the emergence of a nonzero first moment

5.2.1 Identification of the first moment

We begin by quantifying the coherence of the time-dependent field $\psi(\mathbf{x}, t)$ via its temporal power spectrum¹, evaluated at different spatial locations \mathbf{x} . We define the temporal power spectrum of the classical field ψ at position \mathbf{x} , evaluated over a period of length T as

$$H(\mathbf{x}; \Omega) = |\mathfrak{F}^T \{\psi(\mathbf{x}, t)\}|^2, \quad (5.1)$$

where $\mathfrak{F}^T \{f(t)\}$ denotes the Fourier coefficient taken from some arbitrary time origin

$$\mathfrak{F}^T \{f(t)\} \equiv \frac{1}{T} \int_0^T f(t) e^{i\Omega t} dt. \quad (5.2)$$

Here we calculate the power spectrum for a classical field evolving with the trapping and interaction parameters of section 5.1, and energy $E = 12.0 N_c \hbar \omega_r$. Using the Penrose-Onsager (PO) approach introduced in section 3.3.1, we find that the condensate fraction for this classical field is $f_c \equiv n_0 / \sum_k n_k = 0.70$. We choose a sampling period of $40\omega_r^{-1}$, and approximate the integral in equation (5.2) by a discrete sum over 1000 equally spaced samples of the classical field. In practice, we calculate the power spectrum at points in the $z = 0$ plane, and average it over the azimuthal angle in this plane to smooth out fluctuations. We thus obtain the averaged power spectrum as a function of the radius r , which we present in figure 5.1(a). The oscillation frequencies Ω that we measure in the time-dependent field correspond, of course, to energies $\epsilon = \hbar\Omega$ in the quantum mechanical system.

For comparison, on the same figure we also plot the profile of the harmonic trapping potential $V(r)/\hbar$ in this plane (parabolic green (grey) line) and the cutoff energy E_R/\hbar (horizontal line), and the classical turning point (vertical line) of the condensate band defined by their intersection. The most prominent feature in this plot is the strong peak in the power spectrum centred on $\Omega = 11.4\omega_r$, which is a signature of the long-lived temporal phase coherence of the condensate formed in the classical field. We identify the frequency λ_0 where this peak is located as the *condensate frequency*. The broad, lower intensity background spectrum represents the thermally occupied excitations in the classical-field system. It is worth pointing out that, in the central region of the trap ($r \lesssim 4r_0$), the

¹We note that the power spectrum is related to the autocorrelation function of the field by the well-known *Wiener-Khinchin* theorem [215].




Figure 5.1: (a) Power spectrum $H(\mathbf{x}; \Omega)$ of the classical field on the plane $z = 0$ (azimuthally averaged). Green (grey) lines indicate the trapping potential, cutoff energy and corresponding classical turning point of the trap. (b) Space-integrated power spectrum $N(\Omega)$ of the field (see equation (5.5)) as a function of the phase-rotation frequency Ω .

background spectrum is strongly distorted by the presence of the condensate, with positive frequency components extending to approximately $\lambda_0 + E_R/\hbar$, and negative frequency components extending down to approximately $\lambda_0 - E_R/\hbar$ ². This behaviour represents the restructuring of the excitation spectrum of the trap by the interacting condensate, which distorts the single-particle excitations of the system into the familiar Bogoliubov particle-hole pairs (section 2.2.3).

The temporal coherence we observe results from the quasiuniform phase rotation of the condensate: the phase of the condensate exhibits a uniform rotation at frequency λ_0 , superposed with a slow diffusion. The width of the power-spectrum peak here is thus determined by the rate of this global condensate-phase diffusion. On time scales short compared with the characteristic time scale of phase diffusion, the condensate has an approximately constant phase in a frame co-rotating at frequency λ_0 , i.e., short-time averages in this frame yield a nonzero *first moment* $\langle \psi \rangle$ of the classical field. A key observation of this chapter is that time averages constructed in this way are analogous to the *anomalous averages* which arise in symmetry-broken descriptions of Bose condensation [119], where the appearance of nonzero values for expectations of non-gauge-invariant quantities (i.e., the breaking of the phase symmetry) signals the presence of condensation in the field. We thus consider the classical field frequency-shifted by Ω

$$\tilde{\psi}(\mathbf{x}, t; \Omega) = e^{i\Omega t} \psi(\mathbf{x}, t), \quad (5.3)$$

and consider time-averages of this quantity formed from the same set of samples used to construct the power spectrum in figure 5.1(a). We define the time-averaged field

$$\begin{aligned} \phi(\mathbf{x}; \Omega) &\equiv \left\langle \tilde{\psi}(\mathbf{x}, t; \Omega) \right\rangle_t \\ &\left(= \mathcal{F}^T \{ \psi(\mathbf{x}, t) \} \right), \end{aligned} \quad (5.4)$$

where $\langle \dots \rangle_t$ denotes a time average over a given period T ($40\omega_r^{-1}$ in this case). The time-averaged field $\phi(\mathbf{x}; \Omega)$ is therefore the component of the classical field which phase-rotates like $e^{-i\Omega t}$, and its norm square quantifies the total (i.e., space-integrated) power contained in the field at frequency Ω , i.e.,

$$N(\Omega) \equiv \int d\mathbf{x} |\phi(\mathbf{x}; \Omega)|^2 = \int d\mathbf{x} H(\mathbf{x}; \Omega). \quad (5.5)$$

²Note that the cutoff energy refers only to the *single-particle* modes that the field is expanded on. The energies of excitations in the interacting field are raised because of the mean-field potential they experience.

In figure 5.1(b), we plot this power as a function of the frequency Ω , and note that it exhibits a prominent peak at $\Omega = 11.38\omega_r$. We identify the frequency at which the norm square of the time-averaged field (equivalently, the space-integrated power of the field) is maximised as the condensate frequency λ_0 , and the corresponding time-averaged field $\phi(\mathbf{x}; \lambda_0)$ as the classical-field condensate or *mean field*³. A nonzero time-averaged field occurs because the condensate has a reasonably well-defined phase on short time periods. We identify this quasi-definite phase as an analogue of the condensate phase which emerges in symmetry-broken descriptions of Bose-Einstein condensation; in this viewpoint, the first moment $\phi(\mathbf{x}; \lambda_0)$ is the analogue of the condensate wave function $\langle \hat{\Psi}(\mathbf{x}) \rangle$ in such mean-field theories of Bose condensation. For notational convenience, we introduce the norm square $N_0 \equiv N(\lambda_0)$ of the mean field, and the normalised mean-field mode function $\phi_0(\mathbf{x}) \equiv \phi(\mathbf{x}; \lambda_0)/\sqrt{N_0}$. The norm square N_0 corresponds to the *population* of the mean-field condensate mode, and we indeed find $N_0/N_c = 0.706$, in close agreement with the PO value for the condensate fraction ($f_c = 0.70$). To further compare this temporal-coherence method of identifying the condensate with the PO approach, we calculate the overlap of $\phi_0(\mathbf{x})$ with the eigenvector $\chi_0(\mathbf{x})$ obtained by the PO procedure. We find $1 - |\langle \phi_0 | \chi_0 \rangle| \approx 1.4 \times 10^{-4}$, i.e., the condensate orbitals obtained by the two different procedures agree to a very high accuracy.

5.2.2 Temporal coherence and sample length

The results obtained for the mean field have an important dependence on the averaging time. As discussed by Sinatra and Castin, the condensate phase exhibits diffusive evolution with time in the classical microcanonical ensemble [235]. As a consequence, the results of the spectral analysis introduced in section 5.2.1 will depend on the length of time over which we sample the classical field. To gain some insight into the consequences of this phase diffusion, we first consider a simple analytical model of the phase-diffusive condensate mode. We then present and interpret the results we obtain for the classical-field condensate as a function of the time over which we sample the field.

³The identification of the condensate and its eigenfrequency in this manner bears some resemblance to the method of Feit *et al.* [83] for the identification of Schrödinger eigenstates and eigenvalues from a spectral analysis of trajectories of the time-dependent Schrödinger equation. In contrast to that work, however, the condensate ‘eigenmode’ we consider here emerges from the trajectories of a *nonlinear* equation of motion, and can only be obtained from an analysis of the real-time field trajectories (or some other sampling of the PGPE microcanonical density equation (3.24)).

Single phase-diffusive mode

We first consider a single-mode model of the condensate, in which the amplitude $a_0(t)$ of the condensate mode (with condensate frequency λ_0) exhibits phase diffusion. We assume that the mode does not exhibit any number fluctuations, which is precisely the condition that $g_0^{(2)} = \langle |a_0|^4 \rangle_t / \langle |a_0|^2 \rangle_t^2 = 1$, which is well satisfied away from the critical regime associated with the transition to the noncondensed state [23]. We thus have $a_0 = |a_0|e^{i\theta(t)}$ and, defining $\varphi(t) \equiv \theta(t) - \theta(0)$, we assume $\text{var}\{\varphi(t)\} \equiv \langle \varphi(t)^2 \rangle - \langle \varphi(t) \rangle^2 = 2\gamma t$ [235], where γ is the (phase) *diffusion coefficient*, and $\langle \cdots \rangle$ denotes an average over *realisations* of the amplitude $a_0(t)$ (i.e., an ensemble average). This is precisely the behaviour of the *Kubo oscillator* [101] stochastic process, which obeys the (Ito) stochastic differential equation

$$da_0 = \left[(-i\lambda_0 - \gamma)dt + i\sqrt{2\gamma}dW(t) \right] a_0(t), \quad (5.6)$$

where $dW(t)$ is a *real* Wiener increment, which satisfies $\langle dW(t)dW(t') \rangle = \delta(t - t')dt$. By studying this simple model we hope to gain insight into the behaviour of our diffusive condensate mode.

We consider the power spectrum of the mode obtained over a period T ,

$$N^{(0)}(\Omega; T) = \left| \frac{1}{T} \int_0^T dt e^{i\Omega t} a_0(t) \right|^2. \quad (5.7)$$

This power spectrum is itself a stochastic process (developing in T), i.e., it varies between realisations of the oscillator. We therefore consider its mean $\langle N^{(0)}(\Omega; T) \rangle$. Using the known result $\langle a_0(t)a_0^*(s) \rangle = |a_0|^2 \exp[-i\lambda_0(t - s) - \gamma|t - s|]$ [101], we find

$$\begin{aligned} \langle N^{(0)}(\Omega; T) \rangle &= \frac{1}{T^2 (\gamma^2 + \Delta^2)^2} \left\{ \gamma T (\gamma^2 + \Delta^2) \right. \\ &\quad \left. + [e^{-\gamma T} \cos(\Delta T) - 1] (\gamma^2 - \Delta^2) - 2\gamma\Delta \sin(\Delta T) \right\}, \end{aligned} \quad (5.8)$$

which we have written in terms of $\Delta \equiv \Omega - \lambda_0$ for compactness. In the limit of no diffusion ($\gamma \rightarrow 0$) we regain the result $N^{(0)}(\Omega; T) = |a_0|^2 \text{sinc}^2[\frac{1}{2}(\Omega - \lambda_0)T]$ resulting from the convolution of the delta-function spectrum of the deterministic oscillator with the Fourier transform of the boxcar function representing the finite measurement time T [215]. In the limit of a measurement made on a time scale long compared with the characteristic diffusion time (i.e. $\gamma T \gg 1$), we regain the Lorentzian spectrum of the Kubo oscillator $N^{(0)}(\Omega; T) = (2|a_0|^2\gamma/T)/[\gamma^2 + (\Omega - \lambda_0)^2]$.

From equation (5.8), the power measured at the underlying frequency λ_0 of the oscil-




Figure 5.2: Decay of the norm square of the mean field (i.e. the power measured at the condensate frequency) in a single-mode model, as a function of the averaging time. The asymptotic short-time and long-time behaviours of the function are included for comparison.

lator can be obtained by setting $\Delta = 0$, giving

$$\langle N^{(0)}(\lambda_0; T) \rangle = |a_0|^2 \frac{2}{(\gamma T)^2} [\gamma T - (1 - e^{-\gamma T})]. \quad (5.9)$$

For short time periods $T \ll 1/\gamma$ (such as we consider in the main text), the norm square of the mean field decays like $\sim 1 - \gamma T/3$, while at long times it decays like $\sim 2/\gamma T$ (figure 5.2). It is important to note that this same functional form would be exhibited by (e.g.) a complex Ornstein-Uhlenbeck process [101], which one might reasonably assume as a model for a thermally occupied mode [106, 248] in a classical-field approximation. The ‘bare’ (i.e., infinite sampling time) power spectrum of such a mode is similarly Lorentzian, and so we expect the same behaviour for the condensate and the thermal mode, both for two-time correlations ($|\langle a^*(t)a(0) \rangle| \sim e^{-\gamma t}$) and for the measured power $N(\Omega; T)$. The two cases (i.e. condensate and thermal mode) are thus distinguished only by the *time scales* on which the power decays. *Qualitative* differences between the two types of mode only appear in second-order (and higher) correlation functions, which are sensitive to *number* fluctuations.

Multimode description

In general, the condensate mode is only one mode in a multimode field which contains other, thermally occupied modes. We expect the thermal field to be well described in the basis of Bogoliubov modes $\{(u_i, v_i)^T\}$ orthogonal to the condensate mode [47], and thus assume

$$\left\langle \frac{1}{T} \int dt e^{i\lambda_0 t} \psi(\mathbf{x}, t) \right\rangle = \frac{1 - e^{-\gamma T}}{\gamma T} |a_0| \langle e^{i\theta(0)} \rangle \chi_0(\mathbf{x}), \quad (5.10)$$

where, to gain simple insight into our measurements of the field, we assume that the $\{b_j(t)\}$ are complex Ornstein-Uhlenbeck processes which are uncorrelated with one another and with the condensate process. The space-integrated power spectrum of the field is thus

$$\begin{aligned} \langle N(\Omega; T) \rangle &= \langle N^{(0)}(\Omega; T) \rangle \int d\mathbf{x} |\chi_0(\mathbf{x})|^2 \\ &+ \sum_j \langle N^{(j)}(\Omega; T) \rangle \int d\mathbf{x} |u_j(\mathbf{x})|^2 + |v_j(\mathbf{x})|^2, \end{aligned} \quad (5.11)$$

where

$$\langle N^{(j)}(\Omega; T) \rangle = \left| \frac{1}{T} \int_0^T dt e^{i\Omega t} b_j(t) \right|^2, \quad (5.12)$$

behave similarly to $\langle N^{(0)}(\Omega; T) \rangle$, except that they are centred on the frequencies ϵ_j^B/\hbar of the Bogoliubov modes, and attenuate much more rapidly with T ($\gamma_j \gg \gamma_0$). There is therefore power in the field at a range of frequencies, however, for times $T \gg 1/\gamma_j$ we have $\langle N(\Omega; T) \rangle \approx N^{(0)}(\Omega; T) \int d\mathbf{x} |\chi_0(\mathbf{x})|^2$ and, moreover,

$$\left\langle \tilde{\psi}(\mathbf{x}; \lambda_0; T) \right\rangle = \frac{1 - e^{-\gamma T}}{\gamma T} |a_0| \langle e^{i\theta(0)} \rangle \chi_0(\mathbf{x}), \quad (5.13)$$

where the appearance of the expectation $\langle e^{i\theta(0)} \rangle$ of the initial complex phase emphasises that the condensate phase varies randomly between ensemble members, ‘breaking’ the U(1) symmetry in any particular realisation.

Measured results

From the foregoing discussion, we expect the power in the classical field measured at the condensate frequency to decay with time, exhibiting a power-law tail $N(\lambda_0; T) \sim 2/\gamma T$ at long times. We illustrate this issue using the same simulation ($E = 12N_c \hbar \omega_r$) as in the previous section. Increasing the sampling period to $T \gtrsim 50\omega_r^{-1}$ the condensate frequency is more accurately resolved as $\lambda_0 = 11.39\omega_r^{-1}$. We assume this value as a




Figure 5.3: Norm square N_0 of the time-averaged field (i.e., space-integrated power of the classical field measured at the condensate frequency) as a function of the sampling period T . The solid line shows the value obtained from a single contiguous sampling of the classical field over period T . Circles (with error bars) show the mean (and standard deviation) of estimates N_0 measured on 10 individual $400\omega_r^{-1}$ sub-periods of the $4000\omega_r^{-1}$ time series. The dashed line shows the expected (ensemble averaged) power, extrapolated from a least-squares fit to the means of the short-time estimates.

best estimate for the condensate frequency, and calculate the power at this frequency as a function of the measurement period T , up to a maximum measurement period of $4000\omega_r^{-1}$. In figure 5.3 we plot the power measured at frequency λ_0 (solid line), and find that it decays in a nonuniform way with increasing T . However, the (normalised) mean-field orbital $\phi_0(\mathbf{x})$ we obtain at frequency λ_0 satisfies $1 - |\langle\phi_0|\chi_0\rangle| \lesssim 10^{-4}$ for all averaging periods T we consider; i.e., although the measured *occupation* of the condensate decays with increasing averaging period due to the diffusion of the condensate phase, the *mode shape* we obtain is relatively unaffected. The nonuniform decay of the mean-field orbital's occupation we observe is to be expected for a single trajectory, whereas we expect the scaling $N \sim 2/\gamma T$ to emerge from an average over a large ensemble of similarly prepared classical-field trajectories (cf. [235]). It is possible, however, to infer γ from a single trajectory, as we now show. We divide the total $4000\omega_r^{-1}$ (10^5 -sample) period of the classical-field trajectory into 10 consecutive sub-periods of length $400\omega_r^{-1}$ (each of 10^4 samples), and regard these sub-periods as an ensemble of 10 distinct trajectories. For each member of the ensemble we calculate the power $N(\lambda_0; T)$ (using the best estimate $\lambda_0 = 11.39\omega_r^{-1}$) as a function of $T \leq 400\omega_r^{-1}$. We then average over these 10 ensemble

members to obtain a mean power estimate for each sampling period T . The means and standard deviations of these measurements are indicated by circles with error bars in figure 5.3, and by performing a least-squares fit of the expected power $\langle N(\lambda_0; T) \rangle$ at the condensate frequency (equation (5.9)) to these mean power estimates, we estimate the phase-diffusion coefficient $\gamma \approx 10^{-4}\omega_r$ ⁴. The dashed line in figure 5.3 extrapolates the expected behaviour of $\langle N(\lambda_0; T) \rangle$ to later times. Given this decay of the peak power with T , a rigorous estimate of the condensate population would in principle be obtained by forming estimates $\langle N(\lambda_0; T_i) \rangle$ for multiple sampling period lengths T_i , and extrapolating the resulting trend back to $T = 0$ to estimate the ‘true’ condensate population. However, due to the weak linear decay of the power spectrum peak at short sampling periods, we can accurately estimate the condensate population as the magnitude of the dominant peak in the power spectrum obtained over a short sampling period, for all but the smallest condensate fractions, as we discuss in section 5.3.2.

5.3 Dependence of the first moment on the field energy

In the ergodic classical-field (PGPE) method, equilibrium field configurations of different temperatures can be formed simply by varying the (conserved) energy of the random initial field configuration [67]. In each case the ergodicity of the field evolution leads to the appropriate microcanonical thermal equilibrium (section 3.3.2). In this section we investigate the behaviour of the first moment $\phi(\mathbf{x}; \lambda_0)$ introduced in section 5.2.1 as the energy (and thus temperature) of the classical-field equilibrium is varied, and compare its mode shape $\phi_0(\mathbf{x})$ and occupation N_0 with the Penrose-Onsager condensate orbital $\chi_0(\mathbf{x})$ and occupation n_0 , respectively. We further compare the condensate frequency λ_0 to the microcanonical chemical potential μ of the field obtained using the Rugh methodology introduced in section 3.3.2.

5.3.1 Condensate fraction

We consider here the norm square N_0 of the first moment $\phi(\mathbf{x}; \lambda_0)$ defined as in section 5.2.1, for various values of the classical-field energy $E[\psi]$. In figure 5.4(a) we present estimates N_0 for a range of classical-field energies, and compare them with the condensate occupations calculated by the PO approach. The corresponding classical-field tempera-

⁴We calculated (and fitted to) estimates of the powers at $T = 4, 8, \dots, 400\omega_r^{-1}$, but for clarity indicate only every tenth such estimate in figure 5.3.

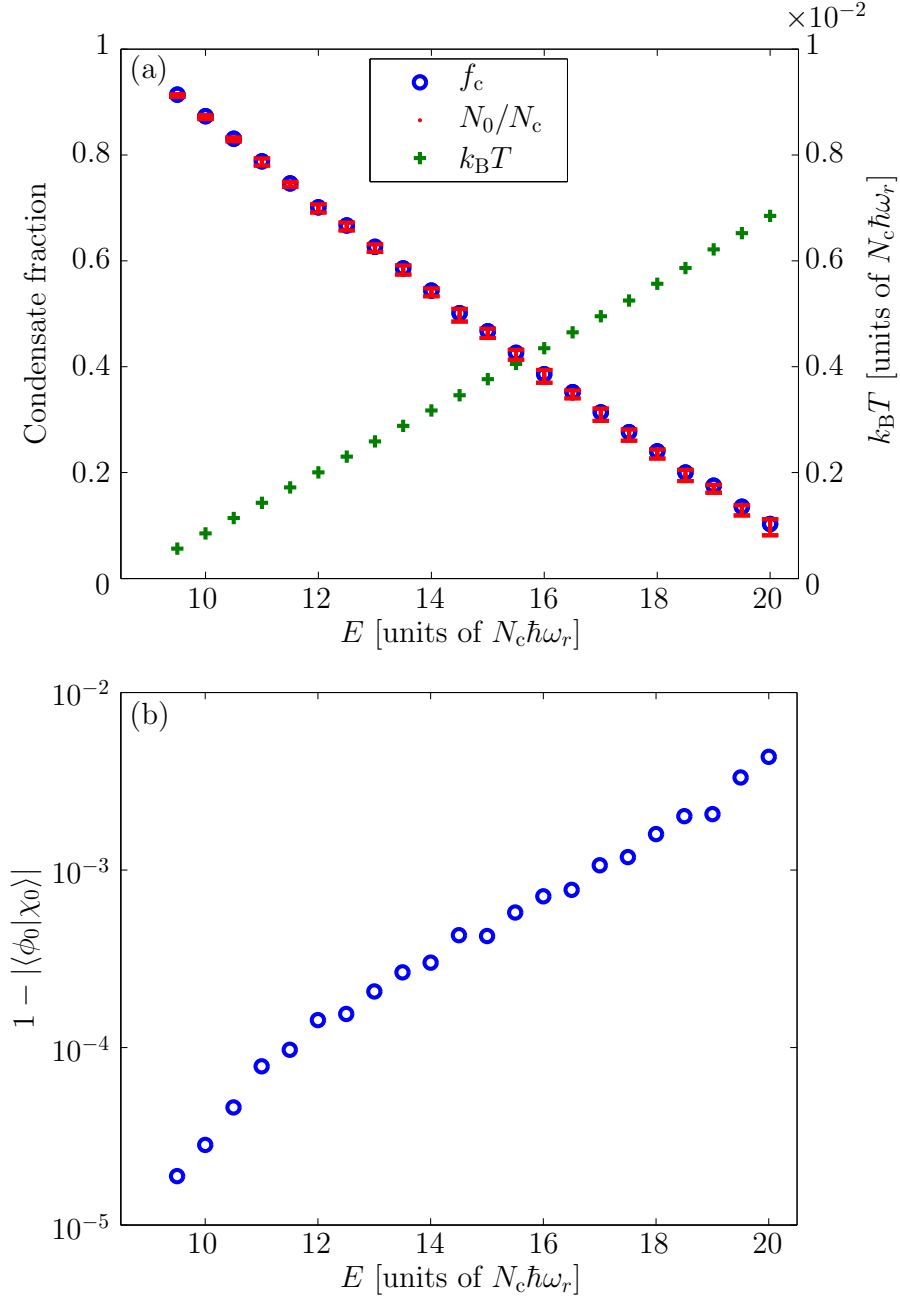


Figure 5.4: (a) Condensed fraction of the classical field, as determined by the Penrose-Onsager procedure (circles) and by the time-averaging procedure (dots with error bars). Plusses indicate the microcanonical temperature of the field. (b) Discrepancy $1 - |\langle \phi_0 | \chi_0 \rangle|$ between the (unit-normalised) first moment $\phi_0(\mathbf{x})$ and the most highly occupied natural orbital $\chi_0(\mathbf{x})$ of the one-body density matrix.

tures, calculated using the Rugh methodology, are also included in the figure. In practice, we calculated the PO condensate by constructing the one-body density matrix (equation (3.26)) from 3000 equally-spaced samples of the classical field taken from a period of $1200\omega_r^{-1}$ of the field evolution. We then divided this period into 30 consecutive sub-periods of length $40\omega_r^{-1}$ which we sampled at a higher resolution (1000 samples per sub-period), from which we obtained 30 separate estimates of the classical-field first moment $\phi(\mathbf{x}; \lambda_0)$. In each sub-period we obtain the mean field as the time-averaged field with the greatest norm (as λ_0 is varied), and we obtain (generally) distinct estimates of N_0 , $\phi_0(\mathbf{x})$ and λ_0 from each series. In this way we exploit the ergodic character of the classical field to emulate sampling from an ensemble of similarly prepared trajectories, as in section 5.2.2. The red (grey) data points in figure 5.4(a) and their error bars represent in each case the mean and standard deviation, respectively, of the norm squares of the 30 estimates of the mean-field. We observe that these estimates agree very closely with the PO condensate fractions f_c (blue circles) throughout the range of energies presented.

We also compare the mean-field orbitals obtained from the time-averaging procedure with the condensate orbitals obtained from the PO approach (see section 3.3.1). In figure 5.4(b) we plot the quantity $1 - |\langle \phi_0 | \chi_0 \rangle|$ (averaged over the 30 estimates) as a measure of the discrepancy between the two orbitals. We observe that for the energies presented ($E \leq 20N_c\hbar\omega_r$) the mean discrepancy is $< 10^{-2}$. At higher energies (corresponding to condensate fractions $f_c < 0.1$), our temporal-coherence approach to identifying the condensate begins to break down: the mean-field orbital $\phi_0(\mathbf{x})$ fails to match the PO orbital $\chi_0(\mathbf{x})$ (i.e. $|\langle \phi_0 | \chi_0 \rangle| < 0.9$) in an increasing fraction of estimates as the condensate fraction $f_c \rightarrow 0$, and so for clarity we have not presented estimates for these energies here. This point is discussed further in section 5.3.2.

5.3.2 Condensate frequency

We consider here the dependence of the condensate frequency λ_0 on the energy of the classical field. By our analogy between the first moment of the classical field and the condensate wave function in mean-field theories (section 5.2.1), we associate this condensate frequency with the *condensate eigenvalue* appearing in such theories, which is itself closely related to the thermodynamic chemical potential of the degenerate Bose-gas system [183]. In figure 5.5(a) we plot estimates of the condensate frequency (red crosses), together with the thermodynamic chemical potential μ (blue circles with connecting line) of the classical field obtained from the Rugh analysis. At the very highest energies, we present results only for ensemble members for which our first-moment analysis and the PO approach agree (i.e. ϕ_0 and χ_0 overlap to within 10%). We observe that the condensate

frequencies λ_0 and the chemical potentials μ agree very well for energies $E \lesssim 20N_c\hbar\omega_r$. Above this energy, the condensate frequencies λ_0 are consistently greater than the chemical potentials. This is expected behaviour, as at a fixed *total* number of system particles, the two quantities differ by a factor of order $1/N_{\text{cond}}$, where N_{cond} is the condensate occupation [218, 183]. Davis *et al.* [68] argued that equipartition of energy in the classical-field model predicts the relationship

$$\mu = \hbar\lambda_0 - \frac{k_B T}{N_0}. \quad (5.14)$$

In figure 5.5(a), we plot the quantity $\lambda_0 - k_B T / \hbar N_0$ (black plusses), where the temperature T is that obtained from the method of Rugh, and find that our results are in reasonable agreement with the prediction of equation (5.14).

We now consider how the total power spectrum of the classical field varies as a function of the field energy. In figure 5.5(b) we plot the power spectrum $N(\Omega)$, averaged over the 30 individual $40\omega_r^{-1}$ sampling periods, for field energies in the range $E \in [9.5, 24]N_c\hbar\omega_r$. At the lowest energies the behaviour of $N(\Omega)$ is as in figure 5.1(b): the function exhibits a prominent peak which we identify with the condensate, and a broad background that we associate with thermal excitations. As the energy (and thus temperature) of the field is increased, the condensate peak decays, and the ‘wing’ of thermal excitations grows until it is of the same magnitude as the condensate peak, and at the highest temperatures only the thermal background remains. This explains why our approach to identifying the condensate begins to fail as the temperature approaches the phase transition: although a temporally coherent condensate may still be present in the field, it becomes increasingly likely that the peak power in any particular estimate of the power spectrum corresponds instead to thermally occupied modes, which eventually swamp the condensate completely.

5.4 Pairing correlations

In the previous section, we have identified that the condensate present in the classical field can be well characterised as the time-average of the appropriately frequency-shifted field. We now show that more general anomalous moments of the field can be obtained from time-averages in the same phase-rotating frame. In this approach, condensation in the classical field is thus accompanied by the appearance of anomalous moments of all orders, in direct analogy to the emergence of general anomalous correlation functions in symmetry-breaking accounts of Bose-Einstein condensation [119, 216].

In terms of the Fock-space decomposition $\hat{\Psi}(\mathbf{x}) = \sum_i \hat{a}_i Y_i(\mathbf{x})$, the emergence of a mean field in the second-quantised formalism is equivalent to the appearance of nonzero first moments $\{\langle \hat{a}_i \rangle\}$. The next-simplest anomalous averages, the quadratic moments $\{\langle \hat{a}_i \hat{a}_j \rangle\}$

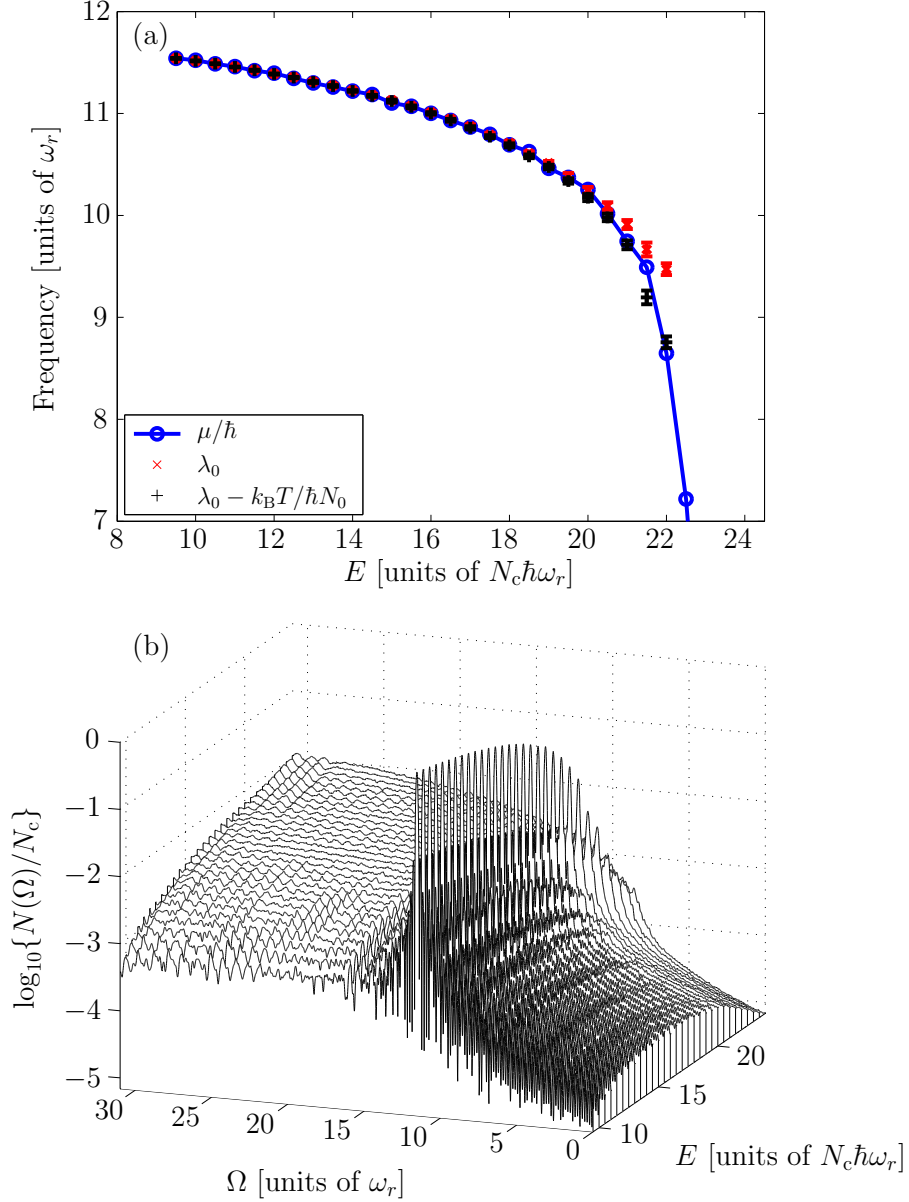


Figure 5.5: (a) Condensate frequency obtained from the time-averaging procedure (red crosses), and microcanonical chemical potential of the classical field (blue circles), for classical-field equilibria with different energies. Black plusses plot the right-hand side of equation (5.14). (b) Norm square of the first moment as a function of the phase-rotation frequency of the frame in which it is constructed (i.e. space-integrated power spectrum of the field), for classical-field simulations with different energies.

(and their conjugates), arise because of the effect of interactions which ‘mix’ the single-particle creation and annihilation operators to form quasiparticle operators $\hat{b} \sim u\hat{a} + v^*\hat{a}^\dagger$ [25]. Consequently, the occupation of quasiparticle modes results in the appearance of nonzero moments of single-particle operators of the form $\langle \hat{a}_i \hat{a}_j \rangle$, which represent correlations between *pairs* of particles. Like the mean field itself, these moments are formally zero in a state of fixed total particle number, although analogous quantities can be defined in particle-conserving terms [183]. Due to the appearance of these pairing correlations, in order to accurately characterise the weakly interacting Bose gas and its excitations, one must consider not only the one-body density matrix $\rho_{ij} = \langle \hat{a}_j^\dagger \hat{a}_i \rangle$, but also the *pair matrix* $\kappa_{ij} = \langle \hat{a}_j \hat{a}_i \rangle$ [25].

In the remainder of this section, we will demonstrate the application of our temporal averaging procedure to the evaluation of quadratic anomalous moments of the classical field. By estimating the pair matrix $\kappa(\mathbf{x}, \mathbf{x}') = \langle \psi(\mathbf{x})\psi(\mathbf{x}') \rangle$ of the noncondensed component of the field, we calculate the *anomalous density*, which characterises pairing correlations in the thermal component of the field. We note that a temporal signature of the anomalous density has previously been observed [40] in homogeneous classical-field simulations, in which the anomalous density is uniform.

5.4.1 Methodology

We seek here to characterise pairing correlations in the thermal component of the classical field, i.e., the component of the field *orthogonal* to the condensate [133], which is obtained by projecting out the condensed component of $\psi(\mathbf{x}, t)$, i.e.

$$\psi^\perp(\mathbf{x}, t) = \psi(\mathbf{x}, t) - \phi_0(\mathbf{x}) \int d\mathbf{x}' \phi_0^*(\mathbf{x}') \psi(\mathbf{x}', t). \quad (5.15)$$

It is important to note that we form $\psi^\perp(\mathbf{x}, t)$ on a given $(40\omega_r^{-1})$ time period by projecting out the mean field obtained over the *same* period, so that (anomalous) averages constructed from $\psi^\perp(\mathbf{x}, t)$ over this period are formed on the same footing as the mean field itself. We transform $\psi^\perp(\mathbf{x}, t)$ to the same phase-rotating frame as the condensate, forming $\tilde{\psi}^\perp(\mathbf{x}, t) = e^{i\lambda_0 t} \psi^\perp(\mathbf{x}, t)$ (where λ_0 is in general different for each period), and then calculate the pair matrix

$$\kappa^\perp(\mathbf{x}, \mathbf{x}') = \left\langle \tilde{\psi}^\perp(\mathbf{x}) \tilde{\psi}^\perp(\mathbf{x}') \right\rangle_t. \quad (5.16)$$

The most well-known characterisation of the anomalous correlations described by the pair matrix is given by the *anomalous density* [119], which we identify as the diagonal part of

the pair matrix:

$$m(\mathbf{x}) = \left\langle \tilde{\psi}^\perp(\mathbf{x}) \tilde{\psi}^\perp(\mathbf{x}) \right\rangle_t \equiv \kappa^\perp(\mathbf{x}, \mathbf{x}). \quad (5.17)$$

We find that the general form of $m(\mathbf{x})$ is apparent from a single estimate of $\kappa^\perp(\mathbf{x}, \mathbf{x}')$, formed over a temporal period $40\omega_r^{-1}$. However, large fluctuations are present in such a single estimate, which is to be expected, as the correlations we seek to resolve here are rather subtle as compared, for example, to the coherence of the condensate. In order to resolve the anomalous density more clearly, we therefore average over multiple estimates of $m(\mathbf{x})$; i.e., for each of 30 consecutive $40\omega_r^{-1}$ periods, we form both the mean field $\phi_0(\mathbf{x})$ and the corresponding anomalous density $m(\mathbf{x})$. The phase of $m(\mathbf{x})$ is only meaningful in relation to the phase of the mean field $\phi_0(\mathbf{x})$ itself, and so for convenience, we choose the overall phase of the classical field in each sampling period such that $\phi_0(\mathbf{x})$ is maximally real. This choice of the phase of $m(\mathbf{x})$ relative to a real and positive condensate wave function corresponds to the traditional choice in mean-field theories. Forming multiple estimates of $m(\mathbf{x})$ in this way allows us to calculate both the mean and the variance of this quantity, as we show in section 5.4.2.

5.4.2 Anomalous density

The mean anomalous density calculated by the procedure described in section 5.4.1 is mostly real and negative (i.e., has phase opposite to that of the mean field), in agreement with the results of mean-field-theory calculations [133, 218, 21], but exhibits some small complex-valued fluctuations due to the finite ensemble size. Denoting averages over estimates by an overbar, we plot in figure 5.6(a) the negative $-\text{Re}\{\overline{m(\mathbf{x})}\}$ of the mean anomalous density on the $z = 0$ plane, and the local standard deviation in estimates $\delta m(\mathbf{x}) = [\overline{|m(\mathbf{x})|^2} - |\overline{m(\mathbf{x})}|^2]^{1/2}$ of the anomalous density on this plane, calculated for a simulation with $E = 14.5N_c\hbar\omega_r$ (for which the condensate fraction $f_c = 0.50$). The anomalous density has the spatial structure expected from mean-field calculations [133, 218, 21]: it resides primarily in the region where the condensate exists, and its absolute value exhibits a shallow ‘dip’ in the centre of the trap. The standard deviation $\delta m(\mathbf{x})$ indicates that the greatest variance in density estimates occurs around the (circular) maximum of $|m(\mathbf{x})|$, while much less variation occurs in estimates of the density in the central dip.

The anomalous density shown here exhibits a very high degree of rotational symmetry about the z axis, but in general the anomalous density we obtain is distorted (the central ‘dip’ in its absolute value becomes saddle shaped along some random axis). We identify this as a result of dipole-oscillation (Kohn mode) [72] excitations of the field, which

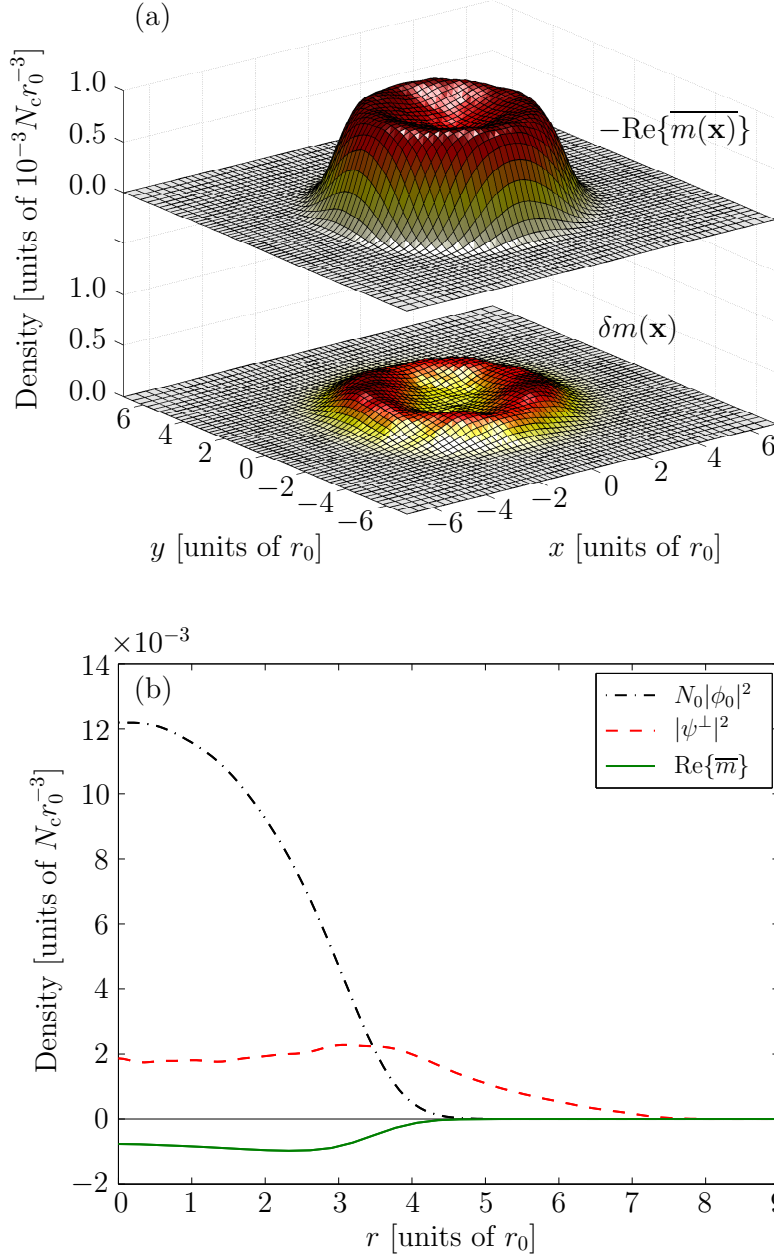


Figure 5.6: Anomalous density $m(\mathbf{x})$ of the field, for the case $E = 14.5N_c\hbar\omega_r$. (a) Shape of $-\text{Re}\{\overline{m(\mathbf{x})}\}$ on a slice through the plane $z = 0$ (upper surface), and standard deviation $\delta m(\mathbf{x})$ of anomalous-density estimates on the same plane (lower surface). (b) Azimuthally averaged density of the condensate mode (as determined by the time averaging), complementary (orthogonal) thermal component of the field, and anomalous density, in the plane $z = 0$.

are immune to thermal damping (section 2.2.1)⁵, and are thus ‘frozen in’ during the thermalisation of the field. More generally, one might regard the classical field as having condensed into an excited centre-of-mass mode, and consider the correlations of the field in a frame following this motion [203]. In figure 5.6(b) we plot the azimuthally averaged anomalous density on the plane $z = 0$, together with the similarly averaged densities of the condensate ($N_0|\phi_0(\mathbf{x})|^2$) and the orthogonal thermal component of the field ($|\psi^\perp(\mathbf{x})|^2$), for comparison. We observe that the magnitude of the anomalous density in the centre of the trap is an appreciable fraction of that of the (normal) thermal component of the field, in agreement with references [133, 218, 21].

5.4.3 Dependence on field energy

Finally, we consider the dependence of the anomalous density on the energy (or equivalently, the temperature) of the projected classical field. In mean-field theories, the anomalous density (after any renormalisation [133, 218, 134]) becomes small as the temperature of the system approaches zero (due to the weak occupation of quasiparticle modes in this limit), and also as it approaches the critical temperature (due to the quasiparticle modes becoming more single-particle-like as the condensate is depleted). This behaviour is often cited as a justification for the neglect of the anomalous density in self-consistent theories (the so-called Popov approximation [119]) in these two limits. In order to characterise more fully the temperature-dependent behaviour of the anomalous density, we follow [133] and calculate its integrated value $M \equiv \int d\mathbf{x} m(\mathbf{x})$. In figure 5.7 we plot the real part of M (neglecting a small imaginary part that arises from incomplete convergence of the averaging; see section 5.4.2) as a function of the classical-field energy. As in our consideration of the condensate eigenvalue in section 5.3, we present for each energy the mean and standard deviation (error bars) of only those estimates obtained from averaging periods which produced an accurate condensate mode (section 5.2.1). We observe that the behaviour of $M(E)$ is consistent with the results of mean-field theories [133, 21], with its absolute value $|M|$ reaching its maximum at intermediate energies (temperatures), and rapidly decreasing as we approach both the zero-temperature and critical regimes.

We note that the well-known issues of ultraviolet divergence of the anomalous density in mean-field theories arise from the *zero-point* occupation of quasiparticle modes [133, 218, 134], which is of course not present in the classical-field model. Also, although the results of classical-field calculations are, in general, dependent on the cutoff energy, the contribution to the anomalous density from successive quasiparticle modes rapidly de-

⁵This property of the many-body system persists in the classical-field model, but note however the effect of the projector on large dipole oscillations [36].




Figure 5.7: Integrated value of the anomalous density $\int d\mathbf{x} m(\mathbf{x})$ as a function of field energy.

creases with increasing energy of the modes, as the modes return to a single-particle structure. The requirement (for our treatment) that the anomalous density is well-contained in the low-energy region (condensate band) described by the PGPE is thus precisely the requirement that the cutoff is effected at such an energy that the interacting Hamiltonian has become approximately diagonal in the single-particle basis $\{\phi_k(\mathbf{x})\}$ (section 3.2.2).

Finally, we note that the maximum (absolute) value of the integrated anomalous density occurs when $f_c \approx 0.5$, but remind the reader that this refers only to the proportion of the *below-cutoff* field that is condensed. Although the entire anomalous density should be well-described by the low-energy Hamiltonian PGPE dynamics, one would have to include the contribution of above-cutoff atoms to the *normal* thermal density of the field [28] in order to draw quantitative comparisons with (e.g.) the mean-field-theory calculations of [133].

5.5 Summary

We have demonstrated that in the Hamiltonian PGPE theory, classical-wave condensation is accompanied by long-range temporal coherence limited only by the slow diffusion of the condensate phase. This gives rise to the appearance of a nonzero first moment of the field, as defined by short-time averages in an appropriate phase-rotating frame. We identified the emergence of this moment with the concept of U(1)-symmetry breaking that is

central to self-consistent mean-field theories. We showed that the mean field obtained by short-time averaging agrees well with the condensate identified by the standard Penrose-Onsager approach, except for close to the critical regime associated with the transition to the normal phase. We found that the condensate eigenfrequency we obtain by this temporal analysis exhibits the behaviour predicted for the condensate eigenvalue in the most sophisticated mean-field approaches [183]; i.e., it agrees closely with the thermodynamic chemical potential at low energies and diverges away from it in inverse proportion to the condensate occupation as the phase transition is approached. By calculating the pair matrix and anomalous density of the noncondensed component of the field, we demonstrated explicitly that time averages in the frame rotating at the condensate frequency allow the calculation of more general anomalous moments. We observed the anomalous density to exhibit the expected behaviour [133], with its magnitude reaching its maximum at intermediate temperatures and decreasing as both the $T = 0$ and critical regimes are approached.

Chapter 6

Equilibrium precession of a single vortex

At zero temperature, an off-axis vortex in a Bose-Einstein condensate orbits (or *precesses*) about the trap axis. Such a state represents the simplest nontrivial example of vortex dynamics in a Bose condensate. The precession of such a vortex can be viewed as the vortex being carried by its own *self-induced* superflow [75], arising due to the effective boundary conditions imposed on the superflow by the inhomogeneity of the condensate orbital density [122]¹. Further complexity is attained in finite-temperature scenarios, in which the vortex may be subject to additional forces resulting from its interaction with the thermal field, as discussed in section 2.3.2. In this chapter we consider a finite-temperature Bose-Einstein condensate in a quasi-two-dimensional trap containing a single precessing vortex. We find that such a configuration arises naturally as an ergodic equilibrium of the projected Gross-Pitaevskii equation, when constrained to a finite conserved angular momentum. In an isotropic trapping potential the condensation of the classical field into an off-axis vortex state breaks the rotational symmetry of the system. We present here a methodology to identify the condensate and the Goldstone mode associated with the broken rotational symmetry in the classical-field model, and we examine the variation in vortex trajectories and thermodynamic parameters of the field as the energy of the microcanonical field simulation is varied.

This chapter is organised as follows: in section 6.1, we discuss our simulation procedure, introducing in turn the parameters of our system, and the details of our numerical approach. In section 6.2 we present results of representative simulations, and characterise the equilibrium trajectory of the vortex. In section 6.3 we discuss the identification of the

¹Corrections to this picture due to the finite extent of the vortex core have been investigated in references [188] and [142].

condensate from the fluctuation statistics of the classical field. In section 6.4 we consider the rotational properties of the condensate and the complementary thermal component of the field. In section 6.5 we discuss the dependence of the system behaviour on the internal energy of the field, and in section 6.6 we consider the dependence of our results on the angular velocity of the rotating frame in which the energy cutoff is effected. In section 6.7 we summarise and present our conclusions.

6.1 Simulation procedure

6.1.1 System

We consider here a system harmonically trapped with a trapping potential isotropic in the xy plane:

$$V(\mathbf{x}) = \frac{m}{2} \left[\omega_r^2 (x^2 + y^2) + \omega_z^2 z^2 \right]. \quad (6.1)$$

In this chapter we will restrict our attention to an effective two-dimensional system. Such systems have been experimentally realised ([247, 55]) by setting the trapping frequency ω_z sufficiently high that no modes are excited in the z direction. A 2D description of our system is thus valid provided that [205]

$$\mu + k_B T \ll \hbar \omega_z, \quad (6.2)$$

where μ and T are the system chemical potential and temperature respectively. Provided that the oscillator length $l_z = \sqrt{\hbar/m\omega_z}$ corresponding to the z confinement greatly exceeds the s -wave scattering length a_s , the scattering is well described by the usual 3D contact potential [205], and we obtain an effective two-dimensional equation in which the collisional interaction parameter is simply rescaled to $U_{2D} = 2\sqrt{2\pi}\hbar^2 a_s/ml_z$ [205]. Restricting the system to 2D permits the dynamics of a *point* vortex interacting with the thermal component of the field to be studied, while removing complexities such as line bending and Kelvin-wave excitations of a vortex filament of finite extent, which provide additional mechanisms for vortex dissipation [151]. In this work we consider ^{23}Na atoms confined in a strongly oblate trap, with trapping frequencies $(\omega_r, \omega_z) = 2\pi \times (10, 2000)$ rad/s. The s -wave scattering length is $a_s = 2.75$ nm and the oscillator length $l_z = 468$ nm. Choosing the cutoff such that $E_R < \hbar \omega_z$, we thus obtain a formally two-dimensional projected field theory, as discussed in section 4.1.1.

In our simulations we will often compare our systems to what we will refer to as our principal *ground state*, which we take as the ground Gross-Pitaevskii (GP) eigenmode [60] of the trapped system in a frame rotating at angular velocity $\Omega_0 = 0.35\omega_r$, with eigenvalue

$(\mu_g)_{\Omega_0} = 10\hbar\omega_r$ in that frame. This state consists of $N_0 = 1.072 \times 10^4$ atoms, and contains a single on-axis vortex carrying unit angular momentum per atom (angular momentum and rotation will always refer to that about the z axis in the remainder of this thesis). In an inertial (nonrotating) frame this state has eigenvalue $\mu_g = 10.35\hbar\omega_r$, and energy $E_g = 7.646 \times 10^4 \hbar\omega_r$. In this chapter we will refer generically to any such inertial frame as a laboratory frame (lab frame).

6.1.2 Microcanonical evolution

Microcanonical formalism

As discussed in section 3.3, the PGPE can be used to evolve a general configuration of the classical field to a thermal equilibrium. In contrast to the simulations of chapter 5, in which the initial field configurations were parameterised solely by their energies $E[\psi]$ (equation (3.8)), in this chapter we also fix the conserved angular momentum

$$L[\psi] \equiv \int d\mathbf{x} \psi^*(\mathbf{x}) \hat{L}_z \psi(\mathbf{x}), \quad (6.3)$$

to a prescribed nonzero value². We therefore expect that the field evolves under the action of the PGPE to a thermal equilibrium consistent with the choices of conserved energy (equation (3.8)) and angular-momentum (equation (6.3)) first integrals. In this chapter we choose $L[\psi] = N[\psi]\hbar = N_0\hbar \equiv L_0$, and thus explore the manifold of classical-field configurations with unit angular momentum per particle. At zero temperature (corresponding to complete condensation in the classical-field model employed here, in which no representation of quantum depletion is included), the corresponding microstate (unique up to a choice of phase) is our principal ground state, containing a singly charged, on-axis vortex, with energy $E[\psi] = E_g$. By choosing initial configurations with $E[\psi] > E_g$ and $L_z[\psi] = L_0$, we investigate the thermal equilibria that result when atoms are thermally excited out of the ground condensate mode, while conserving the total angular momentum.

Choice of frame

In applying our classical-field method, we must make a choice of the rotation rate of the frame in which the PGPE is defined (see equation (4.2)), which we will denote by

²We note that the geometry and projector definition of reference [29] yield a conserved angular momentum about the x axis; however, any nonzero angular momenta in the fields evolved in that work were incidental, and of essentially negligible magnitude, as noted by the authors of that paper.

Ω_p in the remainder of this chapter. The only nontrivial consequence of this choice³ is the precise definition of the projector $\mathcal{P} = \mathcal{P}(E_R; \Omega_p)$, and thus of the set of modes which span the low-energy space \mathbf{C} . This dependence of the projector on Ω_p generalises the familiar concept of an energy cutoff in classical-field theory. As is well known [28], the thermodynamic parameters of the classical field at equilibrium depend both on the conserved first integrals of the system and the imposed cutoff. In the present work, we expect that the equilibrium of the system minimises the free energy

$$F = E - TS - \Omega L, \quad (6.4)$$

where $\Omega = (\partial E / \partial L)_S$ is the thermodynamic angular velocity [157]. The utility of the rotationally invariant Gauss-Laguerre projector (section C.3) used here is that the classical-field evolution is not *dynamically* biased towards a particular rotating frame, i.e., the (generalised) Ehrenfest relation (3.21) for the angular momentum becomes

$$\frac{d\overline{L_z}}{dt} = -\frac{i}{\hbar} \overline{L_z V(\mathbf{x})}. \quad (6.5)$$

This ensures that for the isotropic trapping we consider here, angular momentum is conserved, and further that the field is free to rotate at the angular velocity Ω which minimises equation (6.4), which will in general *not* be equal to the projector angular velocity Ω_p . The relationship between the first integrals (E, L) and the thermodynamic variables (T, Ω) of course depends on Ω_p , as different choices of cutoff yield different Hamiltonian systems. We demonstrate this dependence by explicit calculation in section 6.6, and note that the use of a harmonic basis allows the PGPE theory to be extended to include a mean-field description of above-cutoff atoms, whereby insensitivity to (moderate) changes in the cutoff is regained (see [28]). However, as a first application to the precessing-vortex system, we restrict our attention to a strictly Hamiltonian ‘PGPE system’ [29], with low-energy region \mathbf{C} defined by application of the projector in the frame in which we obtained the vortical ground state (i.e., we take $\Omega_p = \Omega_0 = 0.35\omega_r$), and investigate the behaviour of this particular Hamiltonian system by varying its energy at fixed angular momentum.

Procedure

As discussed in section 3.3, for almost any initial configuration far from equilibrium, we expect the field to approach equilibrium over time [160], due to the ergodicity of

³The appearance of the inertial correction $(\Delta H)_{\text{inertia}} \equiv -\Omega_p L_z$ to the single-particle Hamiltonian (2.2) serves only to ensure that the rotating-frame equations of motion faithfully describe the same dynamics as the fundamental inertial-frame equations.

the system. We thus form our initial states as randomised field configurations under the projector defined by cutoff energy $E_R = 3(\mu_g)_{\Omega_0} = 30\hbar\omega_r$ and rotation frequency $\Omega_p = 0.35\omega_r$, with the same values of normalisation N and angular momentum L as our principal ground state, but with higher energies. In general, forming such states is a highly nontrivial task, due to the nonlinearity of the energy functional. In practice, we employ an algorithm which produces a field configuration with energy E and angular momentum L within prescribed tolerances of the desired values. This algorithm is based on the simpler procedure presented in [28] for the formation of random states *not* subject to the additional angular momentum constraint. Given the ground state $\zeta_g(\mathbf{x})$, and a randomised high-energy state $\zeta_r(\mathbf{x})$, both of normalisation N , we mix the two to form an intermediate configuration

$$\zeta_m(\mathbf{x}) = p_0\zeta_g(\mathbf{x}) + p_1\zeta_r(\mathbf{x}), \quad (6.6)$$

with $p_1 = \sqrt{1 - |p_0|^2}$, and then renormalise this field to N . In general this field will not have $L[\zeta_m] = L_0$. We therefore randomly adjust the amplitudes of pairs of ‘opposite’ angular-momentum modes from the set $\{\alpha_{nl}\}$ which comprise $\zeta_m(\mathbf{x})$, rescaling

$$\alpha_{nl} \rightarrow \sqrt{1 - \xi}\alpha_{nl} \quad (6.7)$$

$$\alpha_{n(-l)} \rightarrow \sqrt{1 + \xi \frac{|\alpha_{nl}|^2}{|\alpha_{n(-l)}|^2}}\alpha_{n(-l)}, \quad (6.8)$$

where n and l are chosen randomly. We typically choose the mixing parameter $\xi = 0.1$, and the factor $|\alpha_{nl}|^2/|\alpha_{n(-l)}|^2$ ensures that the total population of the field is unaltered by this adjustment. We repeat this adjustment for random mode pairs $((n, l), (n, -l))$ under the projector until the angular momentum of the field $L = L_0$ to within the desired tolerance. As a result of these adjustments, the energy of the field generally increases. Thus, after obtaining convergence in L , we adjust the energy-mixing parameter p_0 appropriately, and repeat the angular-momentum adjusting procedure. We repeat this process until both L and E have converged to the desired values to within the desired tolerances. This process is obviously probabilistic, and we cannot be certain *a priori* that the result will converge to an acceptable state in any given time-frame. However, we find that the initial conditions we require for our investigations in this chapter are typically found to within satisfactory tolerances ($\Delta L/L_0 = 10^{-6}$, $\Delta E/E_g = 10^{-4}$) in ~ 5 minutes on a standard desktop computer⁴.

We form initial states in this manner with a range of energies. In each case the central

⁴As a source of ‘random’ initial states for our ergodic simulations this procedure is certainly adequate. However, were we to require any particular *distribution* for the initial states (e.g., for explicit ensemble calculations), we would need to adopt a more sophisticated approach [95].

vortex of the principal ground state remains, however the density profile of each state is severely distorted. As all our states have $L = N_0 \hbar$, the inertial-frame and projector-frame energies are in each case related by a constant shift resulting from the fixed rotational kinetic energy of the field: $(E)_0 = (E)_{\Omega_p} + \hbar \Omega_p N_0$. This shift is of order 5% for the ground-state wave function. In this chapter we will discuss the energies of states in the *lab frame*, and quote them as a multiple of the ground-state energy E_g in that frame.

Having formed the initial states we evolve them in the projector frame for a period of 10^4 trap cycles⁵ (abbreviated cyc), using the adaptive ARK45-IP algorithm (section B.2) with accuracy chosen such that the relative change in field normalisation is $\leq 4 \times 10^{-9}$ per time step taken. By 9×10^3 cyc the states have reached equilibrium, as evidenced by the settling of the distribution of particles in momentum space [67]⁶, and we perform our analysis of steady-state properties on the final 10^3 trap cycles of evolution data.

6.2 Equilibrium behaviour of the field

After the equilibration period, the behaviour of the classical field simulations is qualitatively similar for energies in the range $E \in [1.04, 1.15]E_g$. Each exhibits a central high-density region containing a vortex which is displaced off and precessing about the trap axis. Viewed in the lab frame, the precession is in the same sense as the vortex (i.e. positive rotation sense, corresponding to angular-momentum numbers $m > 0$), and is stochastic in nature, with the vortex position fluctuating about a circular mean orbit, in the presence of strong density fluctuations of the background field. As noted in section 6.1.2, the frequency of the vortex precession is in general not equal to the angular velocity of the projector defining the microcanonical system. The condensate is located entirely in the central bulk (see section 6.3); outside this region the field exhibits the high-energy turbulence characteristic of purely thermal atoms in the classical-field model. In figures 6.1(a-f) we plot classical-field densities spanning a single orbit of the vortex in a simulation with conserved energy $E = 1.10E_g$. The behaviour presented in figure 6.1 is representative of all simulations in which the equilibrium condensate contains a precessing vortex. From a range of simulations we observe that the radius at which the vortex precesses increases as the energy of the field is increased. At higher energies it approaches the edge of the condensate and at $E \approx 1.16E_g$ it is ultimately lost into the violently evolving peripheral region of surface excitations and short-lived phase defects. As the

⁵The term ‘trap cycle’ here refers to the period of oscillations in the radial trapping potential, i.e. $T_{\text{osc}} = 2\pi/\omega_r$.

⁶In practice we average momentum-space densities taken over consecutive 10 cycle periods. Aside from fluctuations close to the origin $\mathbf{k} = 0$ arising from the vortex precession, we observe the densities to settle to their (fluctuating) values well before $t = 9000$ cyc.




Figure 6.1: Classical-field density during a single orbit of the vortex, as viewed in the lab frame. The white circle indicates the vortex position, and + marks the coordinate origin (trap axis). Parameters of the classical field are given in the text.

field energy is decreased, the vortex approaches the trap centre. However, we expect that as the energy of the field is lowered towards that of the ground state, the excitation becomes too weak to reliably generate thermal statistics on the time scales of interest, and we suspect that for low energies the background field ‘seen’ by the vortex may not appear thermal on the time scale of its precession. Indeed in such a low temperature regime the effects of quantum fluctuations neglected in our model will become important. We therefore exclude simulations with energies $E < 1.04E_g$ from our analysis. A detailed discussion of the dependence of system observables on the internal energy of the field is presented in section 6.5.

6.2.1 Vortex precession

In order to characterise the vortex motion, we track the vortex location as observed in the laboratory frame. In practice we recorded vortex locations in the simulation with $E = 1.10E_g$, at a frequency of 25 samples per cycle for a period of 400 cycles, starting at $t = 9000$ cyc. We find that the radial distance of the vortex from the trap centre fluctuates about a mean value of $\overline{r_v} = 2.04r_0$, and that the series of radial distances measured (from the trap centre) has standard deviation $\sigma_{r_v} = 0.09r_0$, which is of the same order as the extent of the vortex core (which is approximately given by the healing length $\eta = 0.20r_0$ [60]⁷). The normalised distribution of vortex-displacement magnitudes measured over the 400 cyc period is displayed in figure 6.2(a). The plotted distribution shows some negative skew but this is not a consistent feature across the simulations performed. To determine temporal characteristics of the vortex trajectory, we define a complex vortex coordinate $z_v(t) \equiv x_v(t) + iy_v(t)$, where the coordinate pair $(x_v(t_j), y_v(t_j))$ specifies the vortex position at time t_j , and we calculate the power spectrum of this time series. As direct periodograms are subject to large variances in the estimates they provide at each frequency [215] we split the 10^4 -sample time series into 10 consecutive segments of equal length and average the power spectra obtained from these shorter series (cf. section 5.2.2)⁸. In figure 6.2(b) we plot the power spectrum of the vortex coordinate z_v over the (angular) frequency range $\omega \in [-12.5, 12.5]\omega_r$ permitted by the sample frequency of our data⁹. The most conspicuous feature of the spectrum is the sharp peak (note the logarithmic scale of figure 6.2(b)) corresponding to the precession of the vortex about

⁷We use the peak density of our principal ground state to estimate the healing length η .

⁸In the signal processing literature this is known as *Bartlett’s method* [126] for estimating the power spectrum.

⁹The Cartesian-grid vortex-tracking algorithm employed here (see [44]) produces a systematic bias in the vortex-position estimate on a sub-grid-point scale (much less than r_0). We have subtracted this dc component of the series $\{z_v\}$ before calculating the power spectrum presented in figure 6.2(b).




Figure 6.2: Quantities characterising the motion of the vortex. (a) Normalised histogram of measured vortex-displacement magnitudes. The vertical dashed line indicates the mean. (b) Power spectrum of the complex vortex coordinate $z_v \equiv x_v + iy_v$, revealing the precession of the vortex as viewed in the laboratory frame.

the trap axis. The next most pronounced feature of the power spectrum is a clearly resolved spike at frequency $\omega = -1\omega_r$. A closer inspection of the power spectrum reveals an accompanying spike at $\omega = 1\omega_r$ (see inset). We identify $\omega = 1\omega_r$ as the universal frequency of the persistent dipole-oscillation mode of the harmonically trapped gas (see section 2.2.1). The frequencies $\pm 1\omega_r$ thus result from a small dipole oscillation of the classical field as a whole. Smaller features can be seen in the wings of the power-spectrum distribution, which we associate with thermal fluctuations of the classical field. From the polar form $z_v(t_j) = |z_v(t_j)|e^{i\theta_v(t_j)}$ of the vortex coordinate, we observe that the vortex phase θ_v does not increase linearly with time, but fluctuates randomly as the vortex moves in the dynamically evolving thermal field. Over time the vortex phase diffuses, so that the rotational symmetry of the system is restored in the ergodic density of classical-field configurations, in direct analogy to the restoration of the phase symmetry of the system by the diffusion of the condensate phase discussed in chapter 5. This is illustrated in figure 6.3, where we plot the distribution of vortex phases θ_v as measured in a frame rotating at frequency $\Omega = 0.2610\omega_r$. This frequency corresponds to the peak presented in figure 6.2(b), and has been chosen because it yields the minimum variation in vortex phase over the 400 cyc period we consider. Histograms of the normalised vortex-phase distribution over the first 50 cycles of this period (black bars) and over the full 400 cyc period (green (grey) bars) are plotted. We observe that at short times the fluctuations in θ_v are small and reasonably Gaussian in nature. Over longer time periods the vortex phase drifts significantly. This diffusion of the vortex phase produces the distinctive Lorentz-




Figure 6.3: Normalised histograms of vortex phases θ_v measured in a rotating frame (see main text) over periods of length 50 cyc and 400 cyc, starting from $t=9000$ cyc.

like shape of the power-spectrum distribution plotted in figure 6.2(b). In section 6.3.2 we discuss the effects of this diffusion on the ‘one-body’ correlations of the classical field obtained from time-averaging the classical-field trajectory, and the implications for the identification and characterisation of the condensate in these simulations.

6.3 Field-covariance analysis

We turn now to the issue of identifying the condensate in the classical-field solutions. In vortex-free¹⁰ equilibrium scenarios the method of condensate definition based on sampling the classical-field correlations discussed in section 3.3.1 has proven successful. In practice, the method is performed in the spectral representation, i.e., we construct the covariance matrix

$$G_{ij} = \langle \alpha_j^* \alpha_i \rangle_t \quad (6.9)$$

from the propagation-basis coefficients (section 3.2.2) of a single classical-field trajectory. The indices i, j each index a full set of quantum numbers which label the basis modes (i.e., n and l in the present case of Gauss-Laguerre basis modes), and $\langle \cdots \rangle_t$ denotes an average over time. As discussed in section 3.3, ergodicity of the field evolution [160] implies that averages taken over sufficiently long times approach averages over the microcanonical

¹⁰Vortices may of course occur in all such simulations; we use the phrase vortex-free here to refer to simulations in which any appearance of vortices is incidental, and sufficiently transient as to not affect the statistics under consideration significantly.

ensemble of field configurations, in which the coherence of a classical-field condensate [57] manifests as a single dominant eigenmode of G_{ij} . However, this method of condensate identification has some significant ambiguities in more general scenarios, in particular when symmetries of the Hamiltonian are broken by the appropriate GP state [51]. For example, the standard PO criterion itself may prove inappropriate in situations in which the condensate exhibits centre-of-mass motion (see section 2.2.1), prompting workers to introduce revised definitions of the density matrix appropriate to such systems, in which the centre-of-mass motion is eliminated [203, 271].

Issues of condensate definition also arise in the presence of vortices. In scenarios in which several vortices exist, breaking the rotational symmetry of an otherwise rotationally invariant many-body Hamiltonian, an uncountably infinite degeneracy of the one-body density matrix appears, even in the idealised Gross-Pitaevskii limit of a zero-temperature system [231]. In this chapter we have chosen one of the simplest scenarios of rotational symmetry breaking in a finite-temperature classical field: the precession of a single vortex. Because the presence of the vortex distinguishes different rotational orientations of the condensate, we find that in order to determine the condensate mode we must construct time averages in an appropriate rotating frame. The fact that the (azimuthal) phase of the vortex diffuses over time in any such uniformly rotating frame means that we must abandon the notion of ergodic reconstruction of the microcanonical density by time averages. However, we expect intuitively [29] that in order to quantify the condensation in the field, the averaging time need only be long enough to distinguish the Gaussian fluctuations of thermal or chaotic modes [116, 177] from the non-Gaussian statistics of a quasi-uniform phase evolution characteristic of a condensed component. We find therefore that we can successfully quantify the condensation of the field from such short-time averages, while exploiting the ergodic character of the evolution over longer time scales. In the remainder of this chapter, we will use the term Penrose-Onsager (PO) procedure generically to refer to the construction of the covariance matrix (density matrix) equation (6.9) by a time-averaging procedure, the precise details of which will be in each case specified in the context.

6.3.1 Lab-frame analysis

We consider a simulation with energy $E = 1.05E_g$, which exhibits a vortex precessing at a frequency $\omega_v \approx 0.23\omega_r$, at a radius $\bar{r}_v \approx 1.3r_0$, which should be contrasted with the extent of the central density bulk $r_b \approx 5r_0$. The vortex in this case is displaced from its axis by significantly more than the extent of its core (healing length $\eta = 0.20r_0$), while the core remains comparatively close to the centre of the region of significant density. We




Figure 6.4: (a-e) Density and phase of the five most highly occupied eigenmodes of the covariance matrix, for case $E = 1.05E_g$. (f) Second-order coherence functions of the covariance-matrix eigenmodes versus mode occupation.

proceed by forming the covariance matrix by averaging over 2501 equally spaced samples of the *lab-frame* representation of the classical-field evolution, between $t = 9000$ and 9100 cyc. We diagonalise the density matrix (6.9) to obtain its eigenvectors, the position-space representations of which we denote by $\chi_i(\mathbf{x})$, and the corresponding occupations n_i (section 3.3.1). In figures 6.4(a-e) we plot the five most highly occupied eigenmodes of G_{ij} , which together contain some 95.4% of the total field population.

The density and phase of the most highly occupied mode ($\chi_0(\mathbf{x})$) is displayed in figure 6.4(a) and contains a singly positively charged *central* vortex (with phase circulation of 2π around its azimuth direction). The second most highly occupied mode ($\chi_1(\mathbf{x})$, figure 6.4(b)) is a rotationally symmetric mode with uniform phase, which forms a pronounced peak at the origin. The third most highly occupied mode ($\chi_2(\mathbf{x})$, figure 6.4(c)) is another vortex mode, with vanishing central density. However, this mode is doubly (positively) charged, with the phase varying by 4π around its azimuth.

We note that these three modes bear a strong resemblance to a central-vortex GP

eigenmode ϕ_0 and the u and v^* spinor components [50]¹¹ of its so-called *anomalous* Bogoliubov excitation (also referred to as the lowest core-localised state by some authors [139]). We recall from section 2.3.3 that this is the lowest-energy excitation of the central single-vortex state in the Bogoliubov approximation, and it has *negative* energy in the laboratory frame, which signals the thermodynamic instability of the vortex state in that frame. The angular-momentum numbers $m_{u,v} = (0, 2)$ of the particle (u) and hole (v) functions result naturally from the transposition of the Bogoliubov pairing ansatz to the $m = 1$ vortex scenario [85, 73]. The collective excitation of the anomalous mode results in the displacement of the vortex off-axis in a direction determined by the phase of the Bogoliubov excitation relative to the condensate (see [223]) in the linear expansion¹²

$$\psi = a_0\phi_0 + bu + b^*v^*. \quad (6.10)$$

In the (zero-temperature) Bogoliubov approximation the energy of the anomalous mode determines the frequency of precession of a displaced vortex under its own induced velocity field, in the appropriate linear-response limit (section 2.3.3).

Returning to our classical-field simulation, we calculate the time-dependent expansion coefficients of the eigenmodes $\chi_i(\mathbf{x})$ over the 100 cyc analysis period, defined as

$$\beta_i(t) \equiv \int d\mathbf{x} \chi_i^*(\mathbf{x})\psi(\mathbf{x}, t). \quad (6.11)$$

The average values $\langle |\beta_i|^2 \rangle_t$ of course yield simply the mode occupations n_i , which are plotted for the 10 most highly occupied modes in figure 6.5(a). From the coefficients $\{\beta_i\}$ we form the classical correlation functions [106, 29]

$$g_i^{(2)} = \frac{\langle |\beta_i|^4 \rangle_t}{(\langle |\beta_i|^2 \rangle_t)^2}, \quad (6.12)$$

which indicate the coherence of the modes over the analysis period. The (local) correlation functions $g^{(n)}$ adopt values of $g^{(n)} = n!$ and $g^{(n)} = 1$ for purely chaotic or thermal fields and purely coherent fields, respectively [116, 29]. The distribution of $g^{(2)}$ values over occupation numbers is shown in figure 6.4(f). We find that the $g^{(2)}$ values calculated for the five most highly occupied modes $\chi_i(\mathbf{x}) : 0 \leq i \leq 4$ are all $\lesssim 1.1$. By contrast, the vast majority of the remaining, less populated modes ($i > 4$) have $g^{(2)} \approx 2$. This

¹¹We adopt here the convention for Bogoliubov-mode component labelling of references [50, 183].

¹²We note that the next most highly occupied eigenmodes, $\chi_3(\mathbf{x})$ and $\chi_4(\mathbf{x})$, have the form of vortex modes with charges $q = +3$ and $q = -1$ respectively. These circulations correspond respectively to those of the u and v^* components of the linear excitation of the GP vortex with angular momentum $m = +2$ relative to the GP state [181].

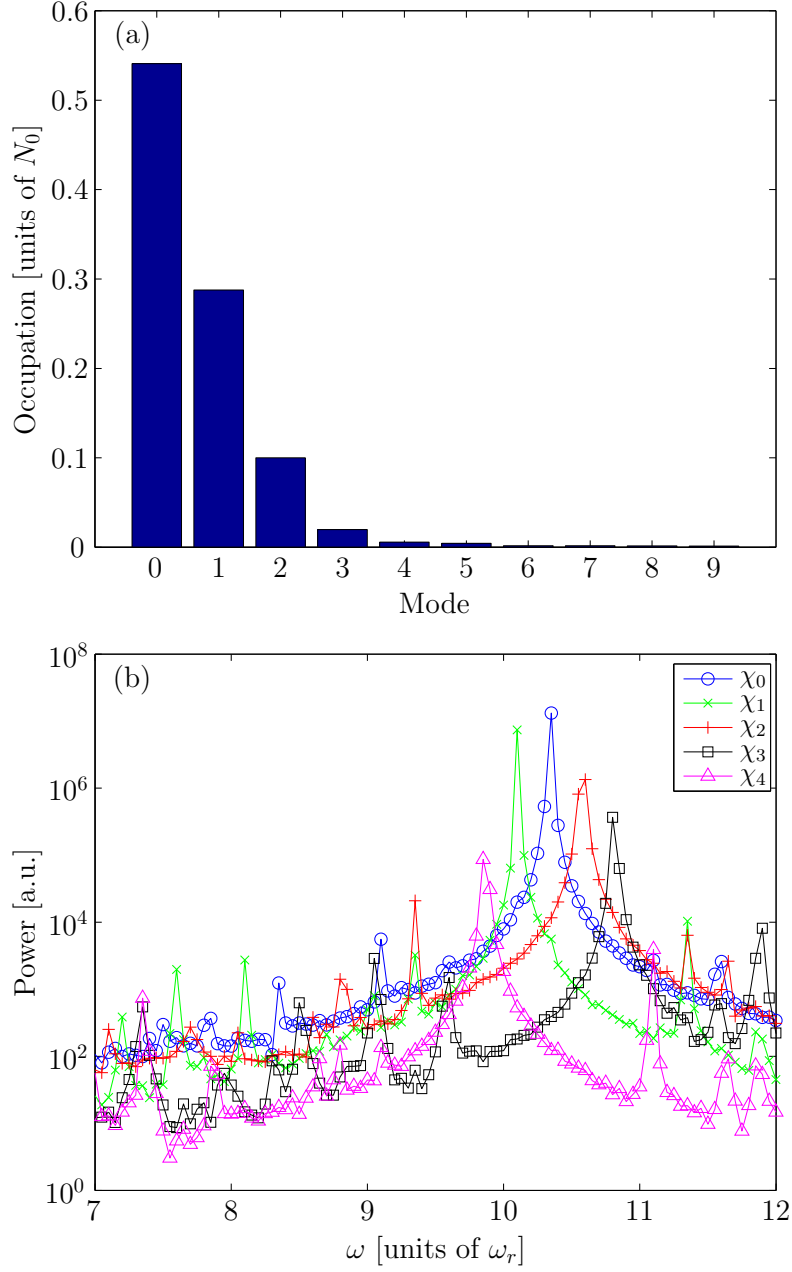


Figure 6.5: (a) Mean occupations of covariance matrix eigenmodes. (b) Temporal power spectra of coefficients $\beta_i(t)$ of the five most highly occupied modes. Quantities presented in each subplot are calculated for the case $E = 1.05E_g$.

suggests that these five most highly occupied modes are *coherent*, while the majority of modes χ_i are to a large degree incoherent *with respect to the averaging procedure employed here*. These coherent modes thus appear to comprise a *fragmented* condensate [164]; i.e., a field configuration in which more than one ‘condensed’ single-particle mode is present. In chapter 5, the temporal power spectrum of the classical field was used to analyse the coherence in the field. Here we apply a similar technique to extract information from the mode coefficients $\beta_i(t)$. In figure 6.5(b) we plot the spectra obtained from these samples, each of which was formed by averaging spectra estimates from five consecutive time series (see section 6.2.1). We observe that the spectrum of each mode coefficient presented in figure 6.5(b) exhibits several peaks, the highest of which is in each case several orders of magnitude greater than the others. We interpret the largest peak in each case as an indication of the quasi-uniform phase rotation associated with coherence in the classical field (see chapter 5). We also note that the frequency peaks of modes χ_1 and χ_2 are spaced evenly about the peak of χ_0 . Specifically we find for the frequencies of peak power $(\omega_{p,0}, \omega_{p,1}, \omega_{p,2}) = (10.35, 10.10, 10.60)\omega_r$, corresponding to a quasiparticle energy of $\varepsilon_q = \hbar(\omega_{p,1} - \omega_{p,0}) \approx -0.25\hbar\omega_r$, in agreement with the vortex precession frequency $\omega_v \approx 0.23\omega_r$ (determined by the technique of section 6.2.1) to within the frequency resolution $\Delta\omega = 0.05\omega_r$.

A precessing-vortex GP state can of course be understood as a linear combination of such angular-momentum eigenmodes. In a frame rotating at the vortex precession frequency (Ω_v), the phases of the various components all rotate at a single common frequency (the condensate eigenvalue μ/\hbar). In the laboratory frame these components have frequencies $\omega_m = \mu/\hbar + \Omega_v m$, and their consequent dephasing over time leads to the apparent fragmentation observed in the lab frame¹³. Simulations of classical fields with higher energies, which contain vortices precessing at larger radii, when subjected to the above analysis, again yield a decomposition into angular-momentum eigenmodes. However, in such cases we find that the most highly occupied mode has $m = 0$, as the anomalous particle component u has grown, pushing the vortex further from the trap centre and replacing the $m = 1$ vortex mode as most highly occupied.

6.3.2 Rotating frames

The above results of a naive application of the PO procedure in the laboratory frame suggest that the classical field condenses in a rotating frame. The mean phase rotation

¹³We note that the smaller frequency peaks in the spectrum of each coefficient correspond to the quasi-uniform phase rotation of other angular-momentum components. This arises due to the imperfect nature of the PO decomposition resulting from (e.g.) the finite sample size used to construct G_{ij} [29].

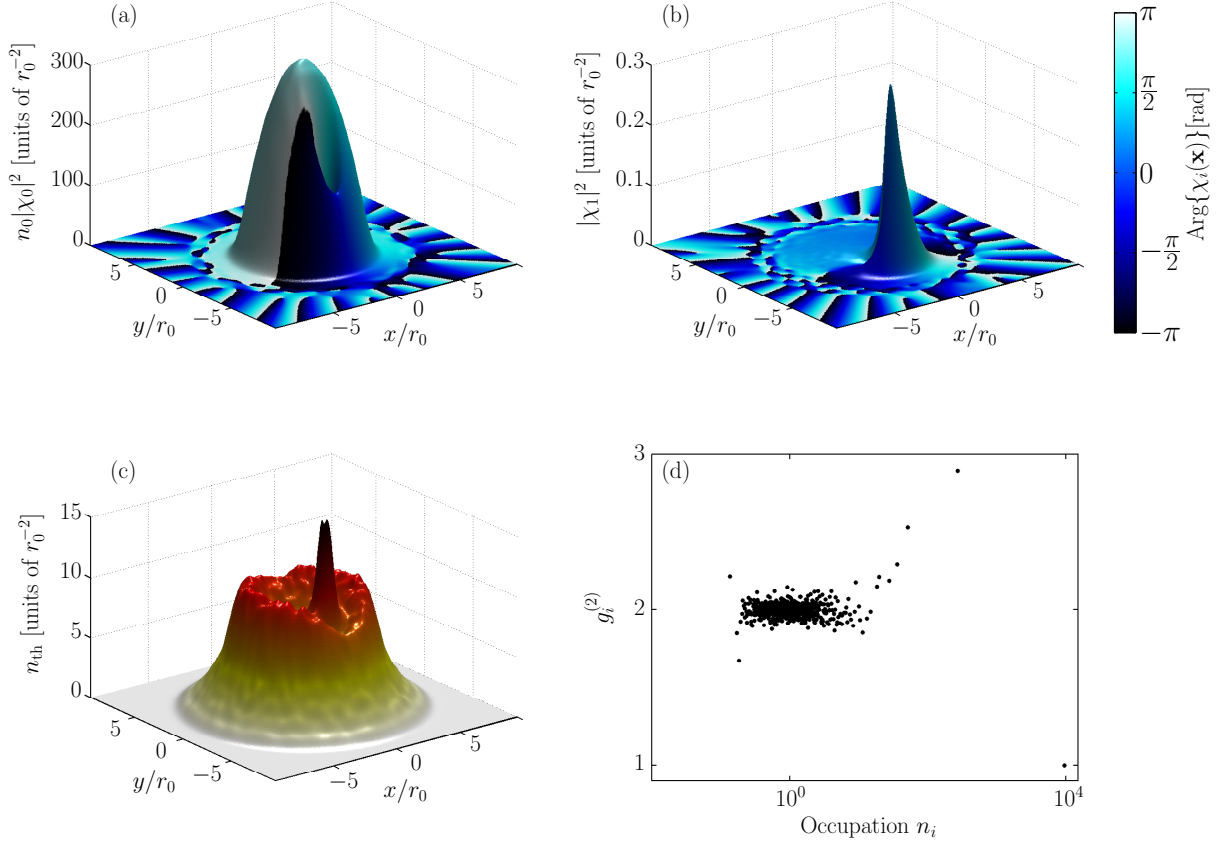


Figure 6.6: Density and phase of (a) the condensate and (b) the Goldstone mode, as determined from the covariance matrix formed in the co-rotating frame. The condensate density is the density of mode $\chi_0(\mathbf{x})$, shown normalised to its occupation n_0 . (c) Thermal density of the field. (d) Second-order coherence functions of the covariance-matrix eigenmodes versus mode occupation.

of the various angular-momentum components is such that upon transformation to an appropriate rotating frame these modes collectively exhibit quasi-uniform phase rotation at a common frequency. We therefore expect that transforming the field to its representation in such a rotating frame, before applying the PO procedure, will yield a single condensate mode. We consider now a simulation with $E = 1.10E_g$, where in equilibrium a single vortex precesses with mean displacement of $\overline{r_v} \approx 2.0r_0$ from the trap centre. After transforming the field coefficients $\{\alpha_{nl}\}$ to a frame (the *measurement* frame) with fixed rotation frequency Ω_m , we construct the covariance matrix in the same manner as described in section 6.3.1. In figures 6.6(a-b) we plot the two most highly occupied modes obtained from this covariance matrix, having performed the PO procedure in a frame with $\Omega_m = 0.2610\omega_r$. We discuss below the reasons for this choice of rotation frequency. The most highly occupied mode is now a vortical state in which the vortex is *displaced*

off-axis, similarly to the case for the full classical field. The second most highly occupied mode is now a peaked function which partially ‘fills’ the density dip associated with the vortex in χ_0 . The prominent peak of this mode exhibits a distinctive amplitude and phase pattern, which we identify as that of the *Goldstone mode* associated with the breaking of rotational symmetry by the off-axis vortex state. Such a mode arises in scenarios in which the condensate mode ϕ_0 breaks a symmetry of the system Hamiltonian (see section 2.2.3). In the case of broken rotational symmetry [169], the Goldstone mode has the form $(u_G, v_G) \sim (L_z \phi_0, -L_z^* \phi_0^*)$, and corresponds to small rotations of the symmetry-broken condensate mode. As such a Goldstone mode is of a fundamentally different nature to the ‘proper’ normal modes (quasiparticles) of the system [25], we regard this mode as separate from the thermal or noncondensate fraction of the field¹⁴. We therefore form the thermal density as the total density of all modes other than the condensate mode (χ_0) and Goldstone mode (χ_1),

$$n_{\text{th}}(\mathbf{x}) = \sum_{i \geq 2} n_i |\chi_i(\mathbf{x})|^2. \quad (6.13)$$

We plot this density in figure 6.6(c). We note that the thermal density exhibits a sharp peak at the vortex location, resulting from the erratic motion of the vortex about its mean trajectory in the classical-field evolution. Thus a careful application of the PO procedure yields a core-filling thermal component, as is predicted by the Hartree-Fock-Bogoliubov family of self-consistent mean-field theories [257]. This is in contrast to the instantaneous classical field, which, due to its essential nature as a single-particle wave function (i.e., vector in the projected single-particle Hilbert space) contains only ‘bare’ vortices, with no possibility of core filling. In the microcanonical-ergodic approach of the PGPE (section 3.3) the two (and higher) particle correlations emerge from the field when averaged over time. However, an alternative interpretation of classical-field trajectories as representative of single experimental trajectories can be made, but, as noted in [28], this interpretation must be made with caution. In particular it is *not correct* to interpret the instantaneous state of the many-body wave function as the Hartree product of the classical field $\psi(\mathbf{x})$ (as is the interpretation of the GP wave function in the zero-temperature GP theory; see section 2.2.2)¹⁵. Rather, the field $\psi(\mathbf{x})$ must be interpreted as an approximation to the quantum *field operator*, from which certain correlations can be deduced. Some averaging of this instantaneous field $\psi(\mathbf{x})$, over an ensemble of trajectories

¹⁴Note that the Goldstone mode only happens to be the second most highly occupied mode in the present case as its spurious ‘population’ is larger than that of all other noncondensate modes.

¹⁵It could indeed be said that this is the definitive distinction of interpretation between the time-dependent GP theory and the time-dependent classical-field theory.

[190] or over time [29], must be performed to construct the correlations necessary to describe the thermal core-filling of a vortex.

We calculate the time-dependent expansion coefficients of the modes $\chi_i(\mathbf{x})$ in the same manner as described in section 6.3.1. The mean occupations of the modes are plotted in figure 6.7(a), and reveal that a single mode contains 89.1% of the field population. We present in figure 6.6(d) the coherence functions obtained from the mode coefficients. The most highly occupied mode is now the only mode with $g^{(2)}$ significantly less than the thermal value $g^{(2)} = 2$. We find in fact $g_0^{(2)} = 1.0002$, indicating very high coherence of the condensate mode over the averaging period. The next most highly occupied modes now have values $g^{(2)} \gtrsim 2.2$, *greater* than that expected for thermally or incoherently occupied single-particle modes. This is a signature of the strongly *collective* nature of the low-lying excitations of the condensate. The form of these excitations is strongly affected by the presence of the interacting, macroscopically occupied condensate, which induces anomalous (pairing) correlations in the field (see section 5.4), which in turn manifest as anomalous values for the coherence functions [103]. We note that the Goldstone mode has $g^{(2)} = 2.90$, close to the upper bound $g^{(2)} = 3$ [103], as we might expect given that a Goldstone mode, with $u_G = -v_G^*$, is the most strongly collective excitation possible. The fact that these pairing correlations are manifest in the one-body density matrix formed in the co-rotating frame is further evidence that this frame gives the correct resolution of the condensate. It is interesting to note that the anomalous nature of these correlations is not apparent in coordinate-space representations of the field [29], or in the lab-frame analysis of section 6.3.1. A more natural characterisation of the low-lying excitations would exploit these anomalous one-body correlations in order to construct an appropriate quasiparticle basis [25].

The choice of measurement-frame frequency ($\Omega_m = 0.2610\omega_r$) optimises the resolution of the condensate as we now discuss. From the discussion of section 6.3.1, we expect (neglecting for now the role of the Goldstone mode) that the classical field can be resolved into a single condensate mode coexisting with a bath of thermal atoms, in some rotating frame. In general the rotation frequency of the measurement frame differs from the equilibrium rotation frequency of the condensate, and this difference leads to dephasing of the distinct angular-momentum components of the condensate mode, and hence to a spurious fragmentation of the condensate into a set of temporally coherent angular-momentum eigenmodes. We propose as a model for analysis of classical-field data, that in the frame in which condensation occurs, the field should decompose into a single condensate mode plus thermal material upon performing the PO procedure. This frame depends on a single parameter (the measurement-frame rotation frequency Ω_m), and the magnitude of the largest eigenvalue, regarded as a function of the measurement-frame

frequency ($n_0 = n_0(\Omega_m)$) forms the corresponding *optimality function*. Our procedure then is to maximise the optimality function to determine the optimal measurement-frame frequency Ω_c , which is the frequency of the rotating frame in which the condensate mode is (most) stationary. This provides a sensitive measure of the (lab-frame) rotation frequency of the vortical condensate mode, and provides a unified basis for the calculation of the properties of the condensed and noncondensed material in the classical field. In particular, we will regard the frame of vortex precession and the frame of condensation as interchangeable, though in practice we will use the latter exclusively in the remainder of this chapter. In figure 6.7(b) we present results of this optimisation procedure as applied to our simulation with $E = 1.10E_g$ over a 100 cyc period starting from $t = 9000$ cyc. We observe a prominent peak in $n_0(\Omega_m)$ at $\Omega_m = 0.2610\omega_r$. We also plot the magnitude of the second largest eigenvalue $n_1(\Omega_m)$, and observe that it exhibits a dip at this same frequency, indicating the sensitivity of the one-body correlations to the frequency Ω_m . We take the location of the peak in $n_0(\Omega_m)$ as an estimate for the rotation frequency Ω_c of the condensate¹⁶.

6.3.3 Temporal decoherence and sample length

We now consider how the condensate mode determined by the above procedure decays as a function of the time over which the averaging procedure is performed. As already noted in section 6.2.1, the phase of the vortex (with respect to any uniform rotation) diffuses over time. This diffusion is intimately related to the breaking of the rotational symmetry of the system by the presence of the off-axis vortex, as we now discuss. The condensate orbital formed in the classical field is an off-axis vortex mode, which is not an eigenstate of angular momentum and thus not rotationally invariant. This breaks the rotational (SO(2)) symmetry of the Hamiltonian equation (3.8)¹⁷. A well-known consequence of the breaking of a continuous symmetry by a dynamical equilibrium in classical mechanics is the appearance of a zero-energy normal mode [117], and this result persists in the quantum theory, where broken symmetries are similarly accompanied by so-called spurious or Goldstone modes [25]. Such excitations are associated with a collective mo-

¹⁶As we vary Ω_m in increments of $\delta\Omega = 10^{-4}\omega_r$, smaller than the frequency resolution afforded by the sampling period ($\Delta\Omega = 1/T_{\text{samp}} = 2\pi \times 0.01\omega_r$), we resolve the profile $n_0(\Omega_m) \sim \text{sinc}^2[(\Omega_m - \Omega_m^0)/\Delta\Omega]$ of the peak of $n_0(\Omega_m)$ (inset to figure 6.7(b)) resulting from its convolution with the boxcar function representing the finite sample time T_{samp} [215]. The location of the central peak of this form provides the best estimate for Ω_c in an effective least-squares sense [95].

¹⁷The rotational symmetry broken here is the invariance of the Hamiltonian under real-space rotations of the field $\psi \rightarrow e^{i\Theta L_z}\psi$. The phase symmetry of the Hamiltonian (invariance under transformations $\psi \rightarrow \psi e^{i\theta}$) is also broken by *any* condensate mode that forms, as discussed in chapter 5. However, the global phase of the classical field has no effect on the density-matrix calculations we perform here.

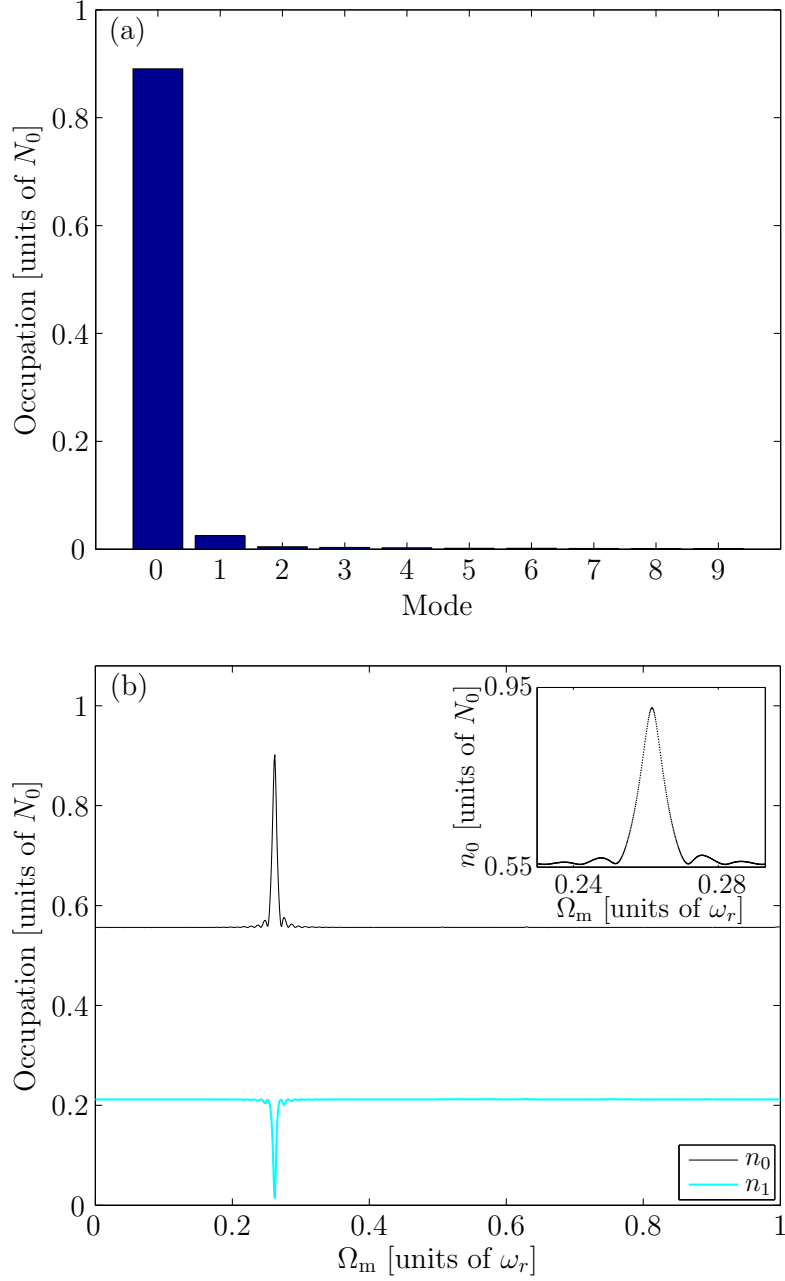


Figure 6.7: (a) Occupations of the 10 most highly occupied eigenmodes of the covariance matrix determined in the frame with $\Omega_m = 0.2610\omega_r$. (b) Dependence of mode populations on the rotation frequency of the frame in which the covariance matrix is constructed. The inset shows the behaviour of n_0 close to its maximum.

tion without restoring force [25, 165]. In the context of classical-field theory, we expect that the erratic evolution of the field results in random excitation of this motion, corresponding to fluctuations of the vortex phase. As there is no ‘restoring force’ associated with this excitation, over time the phase can drift arbitrarily far from any initial value, invalidating the interpretation of the equilibrium as a configuration in which energy is equipartitioned over normal excitations of the symmetry-broken condensate orbital. As a consequence, the decomposition of the field into condensate and noncondensate by the PO procedure also deteriorates over time. This is of course consistent with our expectation that as the averaging time becomes large, the statistics we measure converge to those of the ergodic density of the microcanonical system, which must reflect the rotational symmetry of the Hamiltonian [160], and is entirely analogous to the diffusion of the condensate phase considered in chapter 5, except in this case the Goldstone mode is distinct from the condensate mode itself. In figure 6.8(a) we plot the values of the two largest eigenvalues of the covariance matrix formed by averaging for varying lengths of time, each average starting at $t = 9000$ cycles. For each averaging period we have optimised the condensate fraction over the frequency of frame rotation, which causes only small variation in the optimal frequency ($\lesssim 0.003\hbar\omega_r$). We observe that the largest eigenvalue gets smaller as the averaging period increases, while the second largest eigenvalue gets larger. Over longer periods (inset to figure 6.8(a)) the two eigenvalues approach the values obtained from the density matrix constructed in the laboratory frame (dashed lines), indicating that the orientation of the condensate mode has diffused to such an extent that the density matrix describes a fragmented condensate in all uniformly rotating frames. In figure 6.8(b) we plot the standard deviation of the vortex phases as a function of the averaging period, in a frame rotating at $\Omega = 0.2610\omega_r$. We observe that for the first ≈ 75 trap cycles the standard deviation remains reasonably steady at a small value $\sigma_{\theta_v} \sim 0.05$ radians. After this the standard deviation begins to increase dramatically, and this is ultimately reflected in the measured eigenvalues (figure 6.8(a)). We note that the standard deviation we have presented is simply that of the distribution of vortex phases measured in a single trajectory, as a function of the length of time over which vortex phases are accumulated. To properly characterise the rate of rotational diffusion would require the calculation of the variance of the vortex-phase change over an ensemble of similarly prepared classical-field trajectories (cf. [235]), a numerically heavy task which we do not pursue here. We note further that the initial period of comparative stability of the vortex phase is not a generic feature of the system; in general the excursions of the vortex phase relative to its mean precessional motion occur unpredictably. We note however that the condensate fraction obtained by averaging for only 10 cyc already agrees with the values throughout this period of stability to within 2–3%, in agreement with the

findings of [29]. We therefore adopt the following approach to determining the condensate in the remainder of this chapter: we form the covariance matrix G_{ij} from samples over a period of $T = 10$ cyc, in a frame rotating at frequency Ω_m . We vary this frequency so that the largest eigenvalue of the corresponding covariance matrix is maximised. We take the frequency of this optimal frame as an estimate of the angular velocity of the condensate, and the largest eigenvalue as an estimate of the condensate fraction, and similarly for subsidiary quantities obtained from the PO decomposition (see section 6.4). We do this for 100 consecutive periods of 10 cycles, over the period $t = 9000$ – $10\,000$ cyc and average the estimates for Ω_m and $f_c \equiv n_0 / \sum_k n_k$. In this way we sample the short-time dynamics which determine the correlations of interest, while exploiting the ergodic character of the classical-field evolution over longer times. For the simulation with $E = 1.10E_g$, we find from this procedure $\Omega_c = 0.2598\omega_r$ and $f_c = 0.9110$, in reasonably close agreement with the values obtained from a single averaging over a longer time period. We conclude this discussion by considering the interpretation of the condensate fraction we measure, in light of our abandonment of formal ergodic averaging. We interpret the results of the short-time averaging procedure in terms of a symmetry-broken representation [51] of the many-body system. The one-body density matrix we calculate corresponds to a particular (mean) orientation of the condensate mode (vortex phase θ_v). Labelling this density matrix by G_{θ_v} , the full one-body density matrix is obtained by integrating over all possible vortex phases $G = (1/N) \int d\theta_v G_{\theta_v}$ [51, 231], with N a normalisation factor. This one-body density matrix does not exhibit condensation in the Penrose-Onsager sense, and this is true of the physical N -body system also. Nevertheless, the individual one-body density matrix G_{θ_v} exhibits a condensate, and we interpret the correlations it describes as characterising the behaviour of a corresponding physical system *observed* to have a particular vortex orientation. With this in mind, we will continue to refer to the largest eigenvalue of the density matrix G_{θ_v} constructed from the classical field trajectory as the condensate fraction.

6.4 Rotational properties of the field

A characteristic feature of superfluids is that they resist any attempt to impart a rotation to them, and they acquire angular momentum through the nucleation of vortices¹⁸, which endow the superfluid with *quantised* flow circulation (see section 2.3.2). Thus while a classical fluid in equilibrium with a container rotating at angular velocity Ω satisfies the

¹⁸Condensates formed in dilute Bose gases may possess some angular momentum without containing vortices due to their surface oscillations (see, e.g., [250]).

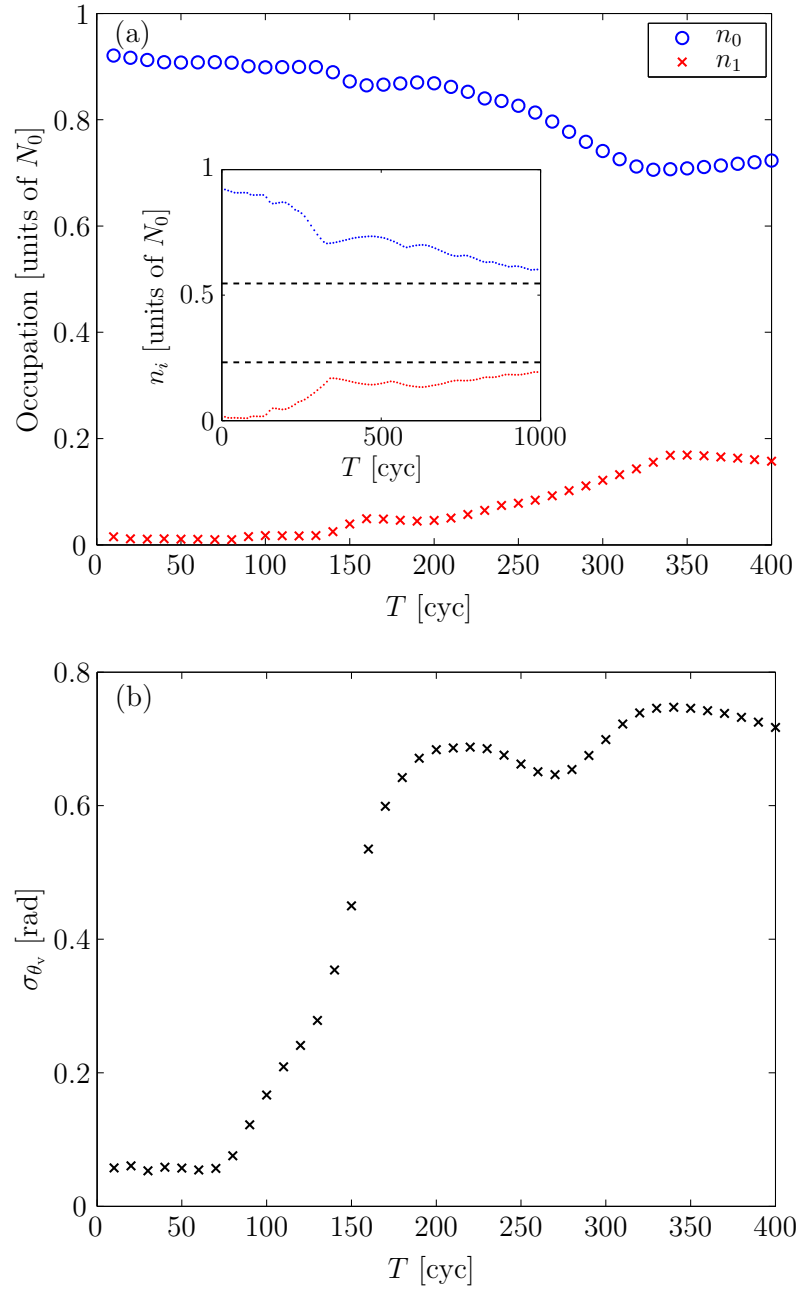


Figure 6.8: (a) Largest eigenvalues (n_0 and n_1) of the covariance matrix as a function of averaging time. The inset shows the behaviour of n_0 and n_1 (upper and lower dotted lines, respectively) over longer times, where the eigenvalues approach their values obtained in the laboratory frame (dashed lines). (b) Standard deviation of vortex phases measured in a frame rotating at $\Omega = 0.2610\omega_r$.

equality

$$\langle L_z \rangle = \Omega \langle \Theta_c \rangle, \quad (6.14)$$

where the classical moment of inertia is defined

$$\Theta_c \equiv mr^2, \quad (6.15)$$

a superfluid in equilibrium with such a container will in general fail to do so. That is to say, the (quantum) moment of inertia

$$\Theta \equiv \frac{\langle L_z \rangle}{\Omega}, \quad (6.16)$$

will not equal the classical value $\langle \Theta_c \rangle$ in such a state. We note that for pure superfluids, such as a dilute Bose gas at $T = 0$ (which is well described as a ground state of the GP equation), the steady-state angular momentum is typically much less than the classical value determined by its mass distribution [273], and this is true also for condensates containing small numbers of vortices [79]. At finite temperatures, we expect the condensed superfluid component of the field to exhibit a suppressed moment of inertia, while the complementary thermal component rotates as a classical fluid.

6.4.1 Condensate and noncondensate observables

The prescription we have developed for decomposing the classical field into condensed and noncondensed parts allows us to extract quantities which characterise the rotation of the two components individually, by calculating appropriate expectation values. We recall from section 2.1.1 that the extension of a single-body operator J to a second-quantised operator $\hat{J} = \int d\mathbf{x} \hat{\psi}^\dagger J \hat{\psi}$ on the many-particle Fock space [88] has expectation value in the many-body state $\langle \hat{J} \rangle = \text{Tr}\{\hat{\rho}_1 J\}$. We thus define an expectation value of a single-body operator in the classical field analogously by

$$\langle J \rangle_c \equiv \text{Tr}\{GJ\}, \quad (6.17)$$

where G is the covariance matrix defined in equation (6.9). The decomposition into a condensed mode and noncondensed modes introduced in section 6.3.2 allows us to write¹⁹

¹⁹In the short-time averages (section 6.3.3) the Goldstone mode is typically not differentiated from the proper noncondensed modes, as its ‘occupation’ is likely to be small. We therefore simply regard all modes orthogonal to the condensate mode $\chi_0(\mathbf{x})$ as comprising the noncondensate field here.

(using Dirac notation for vectors in the projected single-particle space)

$$G = n_0|\chi_0\rangle\langle\chi_0| + \sum_{k>0} n_k|\chi_k\rangle\langle\chi_k| \equiv G_0 + G_{\text{th}}. \quad (6.18)$$

We thus define averages in the condensate and noncondensate by

$$\langle J \rangle_0 = \text{Tr}\{G_0 J\} = n_0 \sum_m \langle m|\chi_0\rangle\langle\chi_0|J|m\rangle, \quad (6.19)$$

and

$$\langle J \rangle_{\text{th}} = \text{Tr}\{G_{\text{th}} J\} = \sum_{k>0} n_k \sum_m \langle m|\chi_k\rangle\langle\chi_k|J|m\rangle, \quad (6.20)$$

respectively. It is clear from the form of G_0 and G_{th} that they are (proportional to) mutually orthogonal projectors, and so we have the additivity property $\langle J \rangle_0 + \langle J \rangle_{\text{th}} = \langle J \rangle_c$.

6.4.2 Angular momentum of the condensate

As noted in section 6.3.2, in the precessing-vortex scenario the condensate determined by our method is a non-axisymmetric state with an off-centre vortex. Analogous condensate orbitals appear in the zero-temperature GP theory as intermediates between the rotationally invariant $\ell \equiv \langle L_z \rangle / N\hbar = 0$ and $\ell = 1$ modes, with fractional angular-momentum (per atom) expectation values $0 < \ell < 1$ [42, 191], and are mechanically unstable (see the discussion in [150]).

Applying the above averaging procedure to our simulation with $E = 1.10E_g$ we obtain $L_c \equiv \langle L_z \rangle_0 / n_0 = 0.63\hbar$. We recall that the condensate fraction obtained for this state was $f_c = 0.911$. Turning to the noncondensate we find $L_{\text{th}} / N_{\text{th}} = 4.80\hbar$. While at this energy, only 9% of the atoms have been excited out of the condensate, in order to maintain rotational equilibrium of the system, the thermal atoms carry nearly half of the angular momentum of the field. This is a consequence of the suppressed moment of inertia of the condensate: in order to have the same angular velocity, the thermal cloud must possess much more angular momentum per particle than the condensate [120, 124].

6.4.3 Angular velocity of the thermal cloud

Now we consider the angular velocity of the thermal cloud. The procedure for separating the condensate and thermal components developed in this chapter allows us analyse the rotation and inertia of both of these two components individually. In particular our approach includes the contribution of thermal fluctuations that traverse the central con-

densed region, as well as those which fill the vortex core, to the total thermal component of the field. To estimate the cloud rotation rate, we assume the expectation value of the classical moment of inertia equation (6.15) evaluated in the thermal cloud as an estimate of the cloud's true moment of inertia, and thus obtain

$$\Omega_{\text{th}} \approx \frac{\langle L_z \rangle_{\text{th}}}{\langle \Theta_c \rangle_{\text{th}}}, \quad (6.21)$$

In this way we find $\Omega_{\text{th}} = 0.262\omega_r$, in fair agreement with the precession frequency of the vortex.

6.5 Dependence on internal energy

We now consider the effect of varying the energy of the classical field on its equilibrium properties. We expect quite generically for Hamiltonian classical-field simulations such as those performed here that an increase in the classical-field energy will result in an increase in the field temperature and a suppression of the condensate fraction [29]. In this scenario with finite conserved angular momentum, we also expect the rotational properties of the system to depend strongly on the field energy.

6.5.1 Density

The most obvious and generic consequence of increasing the energy of the microcanonical system is an increase in the entropy of the system. The qualitative effects of this increase can be readily observed in position-space representations of the classical field density. In figures 6.9(a-f) we plot representative densities of the classical field at equilibrium, for various values of the field energy. These densities are chosen at times in the range $t \in [9000, 10\,000]$ cyc such that the vortex is displaced along the x axis, for ease of comparison. We clearly observe from these images an increase in the surface excitations of the condensate and the development of a turbulent outer cloud as the energy is increased. The images also reveal an increase in the vortex precession radius as the energy is increased.

6.5.2 Condensate fraction and angular momentum

Using the procedure introduced in section 6.3.3, we find the condensate fraction (figure 6.10(a)) drops gradually from $f_c = 0.96$ to $f_c = 0.83$ over the energy range $E \in [1.04, 1.19]E_g$. However, the angular momentum (per particle) of the condensate drops dramatically, falling from $L_c = 0.86\hbar$ at $E = 1.04E_g$ to nearly zero by $E = 1.16E_g$. This




Figure 6.9: (a-f) Representative densities of equilibrium classical fields of different energies. The white circle indicates the vortex position, and + marks the coordinate origin (trap axis).

comparatively severe drop in angular momentum occurs because we are varying the energy of the field at fixed angular momentum; varying the energy at fixed angular velocity we would expect the drop in angular momentum to be commensurate with the depletion of atoms from the condensate.

At higher energies the condensate is vortex-free and thus irrotational, and exists in thermal and diffusive equilibrium with the thermal cloud which now contains all the angular momentum of the field. The small residual angular momentum of these high-energy condensates shown in figure 6.10(a) results from surface excitations which are not differentiated from the condensate by the short-time averaging. We note however that the condensate fraction is essentially continuous across this transition from a symmetry-broken condensate orbital to a non-symmetry-broken (i.e., vortex-free) one, supporting the definition introduced in section 6.3.3 as an appropriate measure of condensation in the symmetry-broken regime.

6.5.3 Condensate and cloud rotation rates

In figure 6.10(b) we plot the condensate rotation frequency, as determined by the frame frequency of optimal condensate resolution (section 6.3.2). We find that the rotation

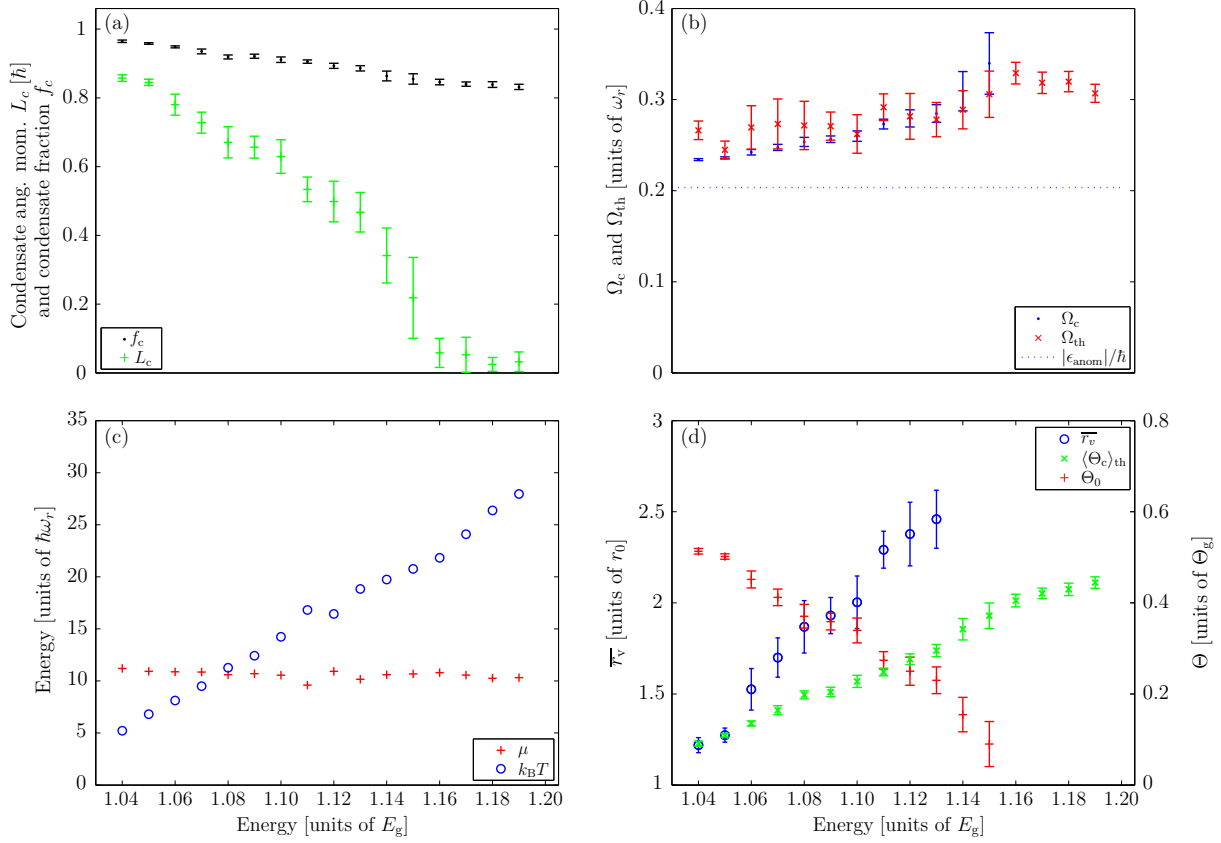


Figure 6.10: Dependence of equilibrium parameters on field energy. (a) Condensate fraction and condensate angular momentum. (b) Angular velocities of the condensate and thermal cloud. (c) Temperature and chemical potential. (d) Vortex precession radius, classical moment of inertia of the thermal cloud, and nonclassical moment of inertia of the condensate mode. Error bars indicate the statistical spread in values calculated (see main text).

frequency increases gradually from $\Omega_c = 0.234\omega_r$ at $E = 1.04E_g$, up to $\Omega_c = 0.340\omega_r$ at $E = 1.15E_g$. For comparison the frequency of the anomalous mode of the ground vortex state in the Bogoliubov approximation, $|\epsilon_{\text{anom}}|/\hbar = 0.2034\omega_r$, is plotted as a dotted line in figure 6.10(b). Note the angular velocity of the condensate increases with *decreasing* condensate angular momentum, in agreement with the well-known thermodynamically anomalous behaviour of the zero-temperature (Gross-Pitaevskii) off-axis vortex state [42, 191, 150]. We note also that the case $E = 1.15E_g$ is very close to the transition from a vortical condensate to a vortex-free one; in $\sim 10\%$ of the sampling periods (each of 10 cyc) the condensate is vortex-free, resulting in maximal values of $n_0(\Omega_m)$ occurring around²⁰ $\Omega_m = 1\omega_r$. We discard these values in calculating the mean and standard deviation presented for this simulation in figure 6.10(b). From figure 6.10(a) we observe that the condensate modes with $E \geq 1.16E_g$ are essentially irrotational, and indeed the condensates in these simulations appear to be largely vortex-free, the majority of samples yielding $\Omega_c \approx 1\omega_r$, with a small minority producing $\Omega_c \sim 0.4\omega_r$, resulting from the transit of short-wavelength surface excitations (so-called *ghost* vortices) which produce a slightly greater calculated condensate fraction in a frame which matches their motion. We thus omit all estimates of the condensate rotation rate for energies $E \geq 1.16E_g$ from figure 6.10(b).

Turning our attention to the cloud rotation frequency (crosses in figure 6.10(b)), determined by the prescription of section 6.4.3, we find these results in qualitative agreement with the measured condensate rotation frequency over the range $E \in [1.04, 1.15]E_g$. The bars accompanying the markers for condensate and cloud rotation rates simply represent the statistical dispersion in the measured quantities as discussed in section 6.3.3, and *not* quantitative measures of the errors in these measurements, but it is clear that the measured cloud rotation rate does not match the condensate rotation rate closely. It is possible that this is associated with pair correlations in the noncondensate, though if anything we expect the presence of pair correlations to suppress the moment of inertia relative to its classical value (cf. [245]). We note finally that the cloud frequency reaches a maximum of $\Omega_{\text{th}} \approx 0.33\omega_r$ at $E = 1.16E_g$ before decreasing as the energy is increased further. This reduction of Ω_{th} with increasing energy is easily understood: once the condensate becomes vortex-free (irrotational), as further atoms are depleted from the condensate the thermal cloud population increases while its angular momentum remains fixed, implying a decrease in angular velocity.

²⁰Where the condensate is vortex-free, we expect that the calculated condensate fraction should be independent of the frame in which the covariance matrix is constructed. However, the condensate fraction often exhibits a small rise around $\Omega_m = 1\omega_r$, which we associate with a persistent dipole excitation of the field.

6.5.4 Temperature and chemical potential

We now estimate the thermodynamic parameters (temperature T and chemical potential μ) of the classical field. Our approach is to construct a fitting function (parameterised by μ and T) and to find the thermal parameters for which it most closely approximates the (time-averaged) density of the classical field. In general, we expect the density of the field to be well described by a semiclassical Hartree-Fock density profile [114], appropriately modified to the classical-field case. In such a semiclassical representation, the rotation of the thermal cloud enters simply as a rotational dilation of the trapping potential experienced by the thermal atoms [251], which then takes the form $V_{\text{rot}}(\mathbf{x}) = m(\omega_r^2 - \Omega_{\text{th}}^2)r^2/2$. However, we would like to avoid performing an iterative calculation to self-consistently solve for the mutually repelling condensed and noncondensed fractions of the field. A simpler approach to constructing the density of the thermal Bose field is presented by the authors of [187], in which the mean-field potential due to the noncondensed atoms is neglected. The thermal cloud is thus constructed as a noninteracting thermal Bose gas which ‘sits’ in the combined potential due to the trap and the mean-field repulsion of the condensate, as in the two-gas model of [74]. However, a straightforward application of this model to our 2D system leads to divergences, as can be seen by considering the density distribution of noncondensed bosonic atoms in the *full* self-consistent semiclassical Hartree-Fock model, which in two dimensions becomes (cf. equation (2) of reference [187])

$$n_{\text{th}}(\mathbf{x}) = \frac{1}{\lambda_{\text{dB}}^2} g_1 \left(e^{-[V_{\text{rot}}(\mathbf{x}) + 2U_{2\text{D}}(n_{\text{c}}(\mathbf{x}) + n_{\text{th}}(\mathbf{x})) - \mu]/k_B T} \right), \quad (6.22)$$

where we have introduced the de Broglie wavelength $\lambda_{\text{dB}} \equiv \sqrt{2\pi\hbar^2/mk_B T}$ [204] and the Bose function $g_\nu(z) \equiv \sum_{j=1}^{\infty} z^j/j^\nu$ [197]. Dropping the mean field potential $2U_{2\text{D}}n_{\text{th}}(\mathbf{x})$ due to noncondensed atoms from equation (6.22) leaves us with a local density approximation for the noncondensed atoms in the effective potential $V_{\text{eff}}(\mathbf{x}) = V_{\text{rot}}(\mathbf{x}) + 2U_{2\text{D}}n_{\text{c}}(\mathbf{x})$. As is well known, the local density approximation for the 2D Bose gas yields a divergent atomic density in the limit $\mu \rightarrow V_{\text{eff}}(\mathbf{x})$, reflecting the absence of thermal-cloud saturation in the homogeneous 2D Bose gas [123], due to the logarithmic divergence of the Bose function $g_1(z)$ as $z \rightarrow 1$ [197].

We are thus led to consider a third approach: we develop a fitting function for the noncondensed density $n_{\text{th}}(\mathbf{x})$ only, which therefore depends only on the *total* field density $n(\mathbf{x})$, which we measure from the classical-field trajectory. The mean-field self-interaction of the noncondensed atoms then ‘lifts’ the bottom of the effective potential above the chemical potential of these atoms, regularising the divergent behaviour and allowing a solution [123]. This has the added benefit that we do not need to distinguish between the

condensed and noncondensed components of the field in order to construct the appropriate mean-field potential²¹.

Our approach to determining μ and T is thus as follows: we assume the noncondensate atoms in the condensate band to be well described by a one-body Wigner distribution, which, for a two-dimensional (Bose) field $\hat{\phi}(\mathbf{x})$ takes the form

$$F(\mathbf{x}, \mathbf{p}) \equiv \frac{1}{h^2} \int d\mathbf{y} \left\langle \hat{\phi}^\dagger \left(\mathbf{x} + \frac{\mathbf{y}}{2} \right) \hat{\phi} \left(\mathbf{x} - \frac{\mathbf{y}}{2} \right) \right\rangle e^{i\mathbf{p} \cdot \mathbf{y}/\hbar}, \quad (6.23)$$

and we further assume for this Wigner function the approximate semiclassical form [47, 16]

$$F(\mathbf{x}, \mathbf{p}) = \frac{1}{h^2} \frac{1}{\exp[(\epsilon(\mathbf{x}, \mathbf{p}) - \mu)/k_B T] - 1}. \quad (6.24)$$

The Wigner distribution is most simply constructed in the frame co-rotating with the thermal cloud (with angular velocity Ω_{th}), where the semiclassical energy is²²

$$\epsilon(\mathbf{x}, \mathbf{P}) = \frac{\mathbf{P}^2}{2m} + V_{\text{eff}}^{\text{HF}}(\mathbf{x}), \quad (6.25)$$

with \mathbf{P} the kinematic momentum in this frame (see appendix D), and the Hartree-Fock effective potential in the frame of the thermal field is thus

$$V_{\text{eff}}^{\text{HF}}(\mathbf{x}) = m \left(\omega_r^2 - \Omega_{\text{th}}^2 \right) \frac{r^2}{2} + 2U_{2D}n(\mathbf{x}), \quad (6.26)$$

with the total field density the sum of the condensate and noncondensate (thermal) contributions $n(\mathbf{x}) = n_c(\mathbf{x}) + n_{\text{th}}(\mathbf{x})$. The factor of two here is the strength of the interaction between pairs of atoms when at least one of the pair is noncondensed [114, 187, 47], i.e., this is the correct mean-field potential experienced by the *thermal* atoms only. The classical-field limit of equation (6.24) is obtained when $F(\mathbf{x}, \mathbf{p}) \gg 1/h^2$, which occurs when the parameter $(\epsilon(\mathbf{x}, \mathbf{p}) - \mu)/k_B T \ll 1$, and is given by its first order expansion in this parameter (cf. equation (3.31)),

$$F_c(\mathbf{x}, \mathbf{p}) = \frac{1}{h^2} \frac{k_B T}{\epsilon(\mathbf{x}, \mathbf{p}) - \mu}. \quad (6.27)$$

It is important to note that although equation (6.27) is obtained here as the high occupation approximation to the bosonic distribution equation (6.24), (describing the classical

²¹This will be of particular importance in chapter 9 where we consider a disordered vortex-liquid state, in which the central quasi-coherent bulk of the field is not ‘condensed’.

²²We assume here that the Hartree-Fock potential is invariant under this rotation, as is the case for the azimuthally averaged potential we construct below.

equipartitioning of energy over phase-space cells of volume h^2), we expect it to apply to the equilibrium state of a weakly interacting classical-field system even if the level occupations are $\lesssim 1$ [239, 170]. The spatial density for the fit is then obtained by integrating equation (6.27) over the momentum \mathbf{p} , up to the *local* momentum cutoff $P_R(\mathbf{x})$ corresponding to the harmonic-basis cutoff defined by the projector \mathcal{P} (see section 4.3.1). This is complicated in the present case by the fact that projector angular velocity Ω_p is distinct from the rotation rate of the thermal cloud Ω_{th} , and so we relegate the details of the calculation to appendix D, where we derive the form of the fitting function

$$n(r; \mu, T) = \frac{1}{\lambda_{\text{dB}}^2} \left[I_1(r; \mu, \Omega_{\text{th}}, \Omega_p) + I_2(r; \mu, \Omega_{\text{th}}, \Omega_p) \right]. \quad (6.28)$$

This expression constitutes an appropriate fitting function for the noncondensed component of the classical field, which requires no explicit knowledge of the distribution of condensed atoms. To obtain an estimate of μ and T for our system, we sample the azimuthally averaged field density $\tilde{n}(r) \equiv \frac{1}{2\pi} \int_0^{2\pi} d\theta n(r, \theta)$ at N_t equally spaced times over some period $t \in [t_0, t_0 + T]$, and average over the sample times to obtain $\langle \tilde{n}(r) \rangle \equiv \sum_{i=1}^{N_t} \tilde{n}(r, t_i) / N_t$. From this we construct the effective potential

$$\langle V_{\text{HF}}^{\text{eff}}(r) \rangle = m (\omega_r^2 - \Omega^2) \frac{r^2}{2} + 2U_{2\text{D}} \langle \tilde{n}(r) \rangle, \quad (6.29)$$

which enters into the fitting function (6.28) through the factors I_1 and I_2 (see appendix D). As our fitting function represents only the noncondensed component of the field, we are obliged to avoid including significant condensate density in our fit to $\langle \tilde{n}(r) \rangle$, and to therefore restrict our fit to a domain $r \in [r_-, r_{\text{tp}}]$, where r_- is the location of the minimum of the effective potential, which we expect to approximately mark the condensate boundary, and $r_{\text{tp}} = \sqrt{2E_R / m(\omega_r^2 - \Omega_p^2)}$ is the semiclassical turning point of the condensate band in the frame of the projector. In practice the low temperatures of the classical-field configurations we consider make performing such a fit difficult; in particular the fit fails close to r_- , in the region of strongly collective excitations of the condensate. We therefore perform our fits only to the wing of the classical-field density; for definiteness we take $r \in [r_- + 1r_0, r_{\text{tp}}]$. We take the values Ω_{th} calculated in section 6.5.3 as estimates for the cloud rotation rate in equation (6.28). Although these values may be somewhat inaccurate as discussed in section 6.5.3, we find that using the vortex precession frequencies (where available) makes little difference to the obtained values of μ and T , and conclude that any discrepancy in Ω_{th} is inconsequential at the level of accuracy of the semiclassical estimates we seek. An example of such a fit is shown in figure 6.11.

The results of this fitting procedure are presented for all field energies E we consider




Figure 6.11: Fitting procedure for the thermodynamic parameters. The solid line illustrates the averaged density of the field times radius. The dashed line indicates the effective potential used in the fit, and the dot-dashed line indicates the fit to the wing of the classical-field density distribution (times radius). The vertical dotted lines indicate the limits of the domain over which the fit is performed.

in figure 6.10(c). We observe that the temperature estimate $k_B T$ increases approximately linearly with the energy E , and is of the same order of magnitude as the chemical potential μ . We note that the chemical-potential estimates are generally higher than the chemical potential of the ground state $\mu_g = 10.35\hbar\omega_r$, whereas we expect the chemical potential of the strictly Hamiltonian PGPE system to decrease with increasing energy (see section 5.3.2). This suggests that the chemical potential is probably overestimated by the fitting, and some error in the temperature estimate is also expected. Nevertheless the results suggest that our simulations probe the temperature regime $k_B T \gtrsim \mu$ in which the thermal friction on the vortex is expected to be appreciable [81], consistent with our observations.

6.5.5 Moment of inertia and vortex displacement

In figure 6.10(d) we plot the classical moment of inertia of the noncondensate (crosses), $\langle\Theta_c\rangle_{\text{th}}$ (see section 6.4.3), in units of the classical moment of inertia $\Theta_g = 7.65 \times 10^4 \hbar/\omega_r$ of the principal ground state. We find that this quantity increases steadily with the field energy as atoms are excited out of the condensate into the thermal cloud. We also plot the (quantum) moment of inertia of the condensate orbital (plusses), which we define as $\Theta_0 \equiv \langle L_z \rangle_0 / \Omega_c$, where the angular momentum $\langle L_z \rangle_0$ and rotation frequency Ω_c of

the condensate are those discussed in the previous sections. To give an indication of the variation of estimates in this derived quantity, we include the standard deviation estimated by adding the standard deviations of $\langle L_z \rangle_0$ and Ω_c in quadrature²³. We note that the moment of inertia of the condensate decays steadily with increasing E , following $\langle L_z \rangle_0$, and by $E = 1.15E_g$ is essentially zero. Above this energy the condensate's angular momentum is basically zero (see the remarks in section 6.5.2), i.e. the condensate fails to respond to the rotation imposed by the thermal cloud, and its moment of inertia is thus zero.

On this plot we also include the displacement of the vortex from the trap centre (circles). For consistency with the other measurements we calculate this by tracking vortex trajectories over 100 successive intervals of 10 cyc (starting from $t = 9000$ cyc) and averaging the measured radii. Bars on the marker again indicate standard deviations of the estimates over the 100 periods. The precession radius increases steadily as the energy is increased. By $E = 1.13E_g$, the precessing vortex is close to the edge of the condensate, and it becomes difficult to track this single vortex in proximity to the proliferation of phase defects undergoing constant creation and annihilation in the condensate periphery. We therefore do not attempt to include data for these higher energies on the plot. It is clear however that the vortex displacement increases as the condensate angular momentum decreases, as is the case in the zero-temperature GP theory [42, 45, 258].

6.6 Dependence on frame of projector

We now consider the effect of the angular velocity Ω_p of the frame in which the projector is defined on the equilibrium configurations of the classical-field trajectories. In irrotational scenarios in which the projector is applied in a laboratory frame, the cutoff energy has a significant effect on the equilibrium attained for given conserved first integrals (energy and normalisation), as discussed in section 3.3.2. In the present case, the cutoff is defined by two parameters: the angular velocity of the frame in which the cutoff is effected, and the (rotating-frame) energy at which the cut is made. As noted in section 6.1.2, the (thermodynamic) angular velocity of the classical field at equilibrium is in general not equal to that of the projector. Here we consider the effect of the choice of projector angular velocity on the classical-field equilibrium. We vary the projector frequency over the range $\Omega_p \in [0, 0.35]\omega_r$, and keep the cutoff energy (in the rotating frame of the projector) fixed at $E_R = 30\hbar\omega_r$. There are several consequences of varying the angular

²³The variations in these quantities are of course not independent; however, the contribution of the variation in estimates of $\langle L_z \rangle_0$ dominates to such an extent that this is of no concern for the indicative representation we seek here.

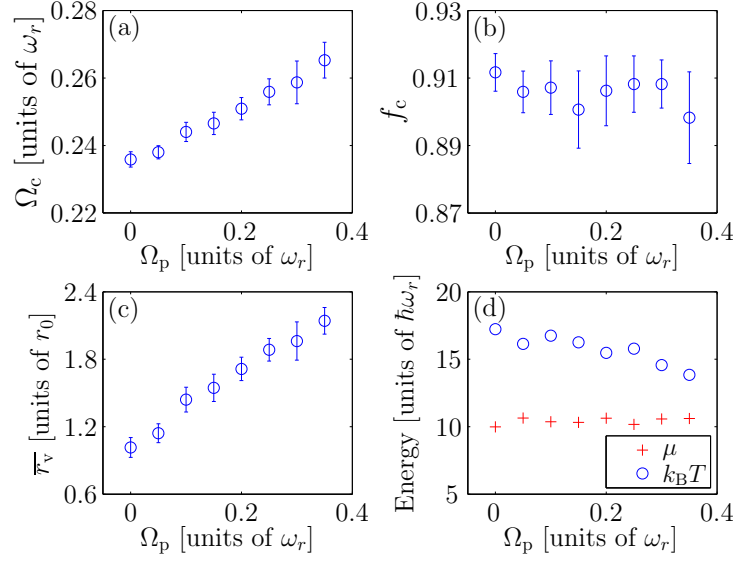


Figure 6.12: Dependence of equilibrium system parameters on the frame in which the projection is effected. (a) Angular velocity of the condensate, (b) condensate fraction, (c) vortex precession radius, (d) temperature and chemical potential.

velocity Ω_p . First, the multiplicity \mathcal{M} (i.e., the number of canonical-coordinate pairs in the Hamiltonian system; see section 3.2.2) depends strongly on Ω_p , increasing with increasing Ω_p (see section C.3.1). Moreover, the modal composition of the coherent region \mathbf{C} varies with Ω_p , with \mathbf{C} becoming increasingly biased towards single-particle modes with $l > 0$ as Ω_p is increased. To illustrate the effect varying Ω_p has on the classical-field equilibria, we have carried out simulations with $E[\psi] = 1.10E_g$, $E_R = 30\hbar\omega_r$, $L[\psi] = N_0\hbar$, and $\Omega_p \in [0, 0.35]\omega_r$. We find that all these simulations exhibit qualitatively similar equilibrium states, each containing a single precessing vortex, as for the simulation with $E = 1.10E_g$ discussed in the main text. We present the key results in figures 6.12(a-d): in figure 6.12(a) we plot the angular rotation frequency of the condensate mode, as determined by the procedure of section 6.3.3. The rotation frequency exhibits a clear upward trend as the frame frequency is increased, nevertheless the dependence of the rotation rate on frame frequency is weak, varying by only $\sim 0.03\omega_r$ over the range of Ω_p considered. In figure 6.12(b) we plot the corresponding condensate fractions. Again the dependence on the frame frequency is weak; any trend in the condensate fraction is of the same magnitude as the statistical dispersion of the measured values. In figure 6.12(c) we plot the vortex precession radii. We see that the radius increases steadily, approximately doubling between $\Omega_p = 0$ and $\Omega_p = 0.35\omega_r$. This is the only quantity which changes significantly over the range of projector angular velocities Ω_p we consider, and as in figures 6.10(b,d) we find that a small proportional change in precession frequency Ω_c

corresponds to a much greater proportional change in $\overline{r_v}$. In fact for the simulations in figures 6.12(a, c) with the smallest precession frequencies, the precession radius $\overline{r_v}$ is both smaller and more strongly dependent on ω_v than for the condensates in figures 6.10(b, d) with comparable vortex precession frequencies, which have greater condensate fractions. This is in agreement with the results of the Gross-Pitaevskii equation for zero-temperature precessing single-vortex configurations, which show that the frequency of vortex precession in less strongly interacting condensates varies less severely with the precession radius of the vortex [191]²⁴.

In figure 6.12(d) we plot the dependence of the chemical potential (plusses) and temperature (circles) on Ω_p , calculated assuming the condensate angular frequency Ω_c as an estimate for the cloud rotation rate. We observe that the temperature decreases as the frame frequency is increased, consistent with the increase in the number of modes in the condensate band (see discussion of equipartition in section 3.3.2). We conclude then that increasing the frame angular velocity Ω_p results in a lowering of the equilibrium temperature, and biases the angular-momentum balance in the system towards the thermal cloud, leading to a higher equilibrium angular velocity. We emphasise that this dependence is a natural consequence of the projector definition: different values of Ω_p define different Hamiltonian systems. Ultimately, this dependence would be removed by the inclusion of above-cutoff effects [104, 65, 28].

6.7 Summary

We have carried out classical-field simulations of a precessing vortex in a finite-temperature Bose-Einstein condensate. Employing a microcanonical ergodic-evolution approach, we find that randomised states of the classical field constrained to have unit (conserved) angular momentum per atom settle to configurations each containing at most a single vortex precessing under the influence of its own induced velocity field, at thermal and rotational equilibrium with the thermal cloud, so that the mutual friction force between the vortex and the cloud vanishes.

As the classical field condenses into a rotationally symmetry-broken state, a naive application of the Penrose-Onsager test of one-body coherence in an inertial frame yields a fragmented condensate comprised of a set of coherent angular-momentum eigenmodes. By transforming the classical field to an appropriate rotating frame, we eliminate the differential phase rotations between the angular momentum components which cause the apparent fragmentation. In this frame we find that the classical field describes a conden-

²⁴Indeed in the ‘lowest Landau level’ limit of very weakly interacting condensates the vortex precession frequency is *independent* of the precession radius [191, 258].

sate mode with an off-axis vortex coexisting with a core-filling thermal component, and we also observe the Goldstone mode associated with the rotational symmetry breaking.

Due to the ergodic nature of the classical-field evolution, the rotational phase of the vortex diffuses over time. Thus care must be taken to separate the short-time fluctuation dynamics which define the coherence of the field from the dynamics on much longer time scales, over which ergodic migration of the field configuration occurs. We found that extracting the condensate by averaging over short times and averaging over successive estimates so obtained allowed us to access the short-time dynamics of interest while exploiting the ergodic nature of the system on longer time scales to better estimate system observables.

We showed that the angular momentum of the condensed mode decreases dramatically as the energy (and thus temperature) are increased at fixed total angular momentum, and correspondingly the vortex precessional radius and frequency increase until the vortex leaves the condensate, and the condensate becomes irrotational. We proposed a measure of the cloud angular velocity based on the decomposition of the field into condensed and noncondensed components and found qualitative agreement with the vortex precession frequency.

Chapter 7

Nonequilibrium dynamics of vortex arrest

In this chapter, we apply the PGPE formalism to a scenario of *nonequilibrium* dynamics of a finite-temperature atomic field. We begin with a finite-temperature precessing-vortex configuration in a rotationally invariant trap, such as those considered in chapter 6. We then introduce an elliptical deformation of the trap which breaks this rotational symmetry, leading to the loss of angular momentum from the atomic field and thus the slowing of the rotating cloud, and consequently the decay of the vortex. Our simulations are the first to describe the arrest of the rotation of both the condensed and noncondensed components of the field in response to the trap anisotropy, and our method has an intrinsic description of the coupled relaxation dynamics of the two components. Adapting the methodology developed in chapter 6, we identify the condensate and the complementary thermal component of the nonequilibrium field, and compare the evolution of their angular momenta and angular velocities. We observe novel behaviour in the relaxation dynamics, arising from the thermodynamically anomalous behaviour of a single-vortex state and its coupling to the nonstationary thermal cloud. By varying the trap anisotropy we alter the relative efficiencies of the vortex-cloud and cloud-trap coupling. For strong trap anisotropies the angular momentum of the thermal cloud may be entirely depleted before the vortex begins to decay. For weak trap anisotropies, the thermal cloud exhibits a long-lived steady state in which it rotates at an intermediate angular velocity.

This chapter is organised as follows. In section 7.1 we describe the initial equilibrium states we consider, and our approach to describing their subsequent evolution. In section 7.2 we analyse the dynamics of a representative classical-field simulation. In section 7.3 we discuss the dependence of the dynamics of the vortex and thermal field on the trap anisotropy, and in section 7.4 we summarise our findings and present our conclusions.

7.1 Simulation procedure

7.1.1 Initial state

As in chapter 6, we choose physical parameters corresponding to ^{23}Na atoms confined in a strongly oblate trap, with trapping frequencies $(\omega_r, \omega_z) = 2\pi \times (10, 2000)$ rad/s. We again adopt a 2D representation of the system, formalised in our classical-field method by choosing a cutoff $E_R = 30\hbar\omega_r \ll \hbar\omega_z$, such that the low-energy space \mathbf{C} excludes all modes with excitation along the z axis (see section 4.1.1). In all our simulations this energy cutoff (projection operation) is effected in an inertial (laboratory) frame, and the low-energy space consists of 465 single-particle propagation-basis modes (see section 3.2.2).

We form a finite-temperature initial state following the procedure of chapter 6: we construct a (nonequilibrium) randomised classical-field configuration over the modes comprising the space \mathbf{C} , with chosen normalisation, energy and angular-momentum first integrals, and we evolve this state for some time (10^4 trap cycles), so that the field has time to migrate to an equilibrium configuration. As in chapter 6, we take as the starting point for forming these configurations the ground Gross-Pitaevskii (GP) eigenstate with $N_0 = 1.072 \times 10^4$ atoms in a frame rotating at angular velocity $\Omega_0 = 0.35\omega_r$, with chemical potential $\mu_g = 10.35\hbar\omega_r$ and energy $E_g = 7.646 \times 10^4 \hbar\omega_r$ in an inertial frame, and angular momentum $L_0 = \hbar N_0$. To this state we add energy and angular momentum, using the algorithm described in section 6.1.2, forming a configuration with $E = 1.10E_g$, and $L = 1.20L_0$. We find that the corresponding equilibrium state is one in which a single vortex precesses very close to the trap axis, i.e., with precession radius $r_v \lesssim \eta$, where the healing length $\eta = 0.20r_0$ is estimated from the density of the ground state (see section 6.2.1). In this configuration the vortex is at equilibrium with the thermal component of the field, and the angular velocity Ω_c of the condensate (which is the same as the cloud's at equilibrium) is close to that at the $\ell \equiv \langle L_z \rangle / \hbar N \rightarrow 1^-$ limit of the mechanically unstable [42, 150] precessing-vortex branch [191]. If additional angular momentum is added to this state, the condensate mode becomes unstable to the nucleation of a second vortex, i.e., the single-vortex state we consider is essentially saturated with angular momentum for the given values of the other conserved integrals (normalisation and energy). We find for this equilibrium initial state $\Omega_c \approx 0.23\omega_r$, $\mu \approx 12\hbar\omega_r$, $k_B T \approx 14\hbar\omega_r$, and condensate fraction $f_c \approx 0.91$.

7.1.2 Evolution

Having formed this initial state, we evolve field trajectories in the laboratory frame in the presence of static trapping-potential anisotropies

$$\delta V(\mathbf{x}) = \epsilon m \omega_r^2 (y^2 - x^2), \quad (7.1)$$

with ellipticities in the range $0.005 \leq \epsilon \leq 0.1$. This procedure corresponds in each case to the sudden introduction of the anisotropy at time $t = 0$. We evolve each of these trajectories with the ARK45-IP algorithm (section B.2), with accuracy chosen such that the relative change in field normalisation is $\leq 10^{-9}$ per time step taken, until the vortex is expelled from the condensate and the total angular momentum of the field has been lost through its interaction with the anisotropy.

7.2 Arrest dynamics

In this section we consider the results of a classical-field simulation with a particular choice of the trap anisotropy, $\epsilon = 0.025$. The relaxation dynamics of the classical field in this case exhibit many features of interest, and serve as a point of comparison for simulations with weaker or stronger anisotropies. Position-space densities of the classical field at representative times are presented in figures 7.1(a-f). Initially the vortex precesses very close to the trap axis (figure 7.1(a)), and it remains close to the trap axis for some time. Figure 7.1(b) shows the density at time $t = 500$ cyc, in which the vortex core remains within $\sim \eta$ of the trap axis. By $t \approx 1000$ cyc the vortex has begun to spiral out of the central density bulk which contains the condensate. Figure 7.1(c) shows that by $t = 1500$ cyc the vortex has undergone significant radial displacement. This increase in radial displacement continues (figure 7.1(d)), until the vortex core approaches the violently evolving condensate boundary (figure 7.1(e)). At $t \approx 3070$ cyc the vortex is lost into the peripheral thermal material, leaving the condensate vortex-free and essentially irrotational (figure 7.1(f)).

7.2.1 Vortex trajectory

We now quantify the behaviour of the vortex during its decay. We observe that the vortex executes an essentially spiral-like motion as it decays, in agreement with the predictions of [81] and the simulations of [227, 140]. However, the vortex trajectory here is strongly stochastic, and the vortex orbits the trap axis many (~ 700) times during its decay, so we do not present the full vortex trajectory. To characterise the radial drift of the vortex, we




Figure 7.1: (a-f) Classical-field densities at various times during the system evolution. The white circle indicates the vortex position, and + marks the coordinate origin (trap axis). Parameters of the simulation are given in the text.

track its location at a frequency of 25 samples per trap cycle, and average the resulting series over intervals of 100 samples (4 cyc) in order to smooth out rapid fluctuations due to thermal density fluctuations in the background field against which the vortex moves and uncertainties introduced by sampling the vortex location on a Cartesian grid. In figure 7.2(a) we present the smoothed vortex radius ($\overline{r_v}$) data. It is clear from figure 7.2(a) that there is an initial lag of ~ 1000 cyc between the introduction of the trap anisotropy at $t = 0$ cyc and the beginning of the upward trend in vortex displacement. We note that the displacement radius does not increase monotonically, but exhibits large oscillations during its increase. The period of these oscillations is often $\gtrsim 100$ trap cycles, spanning many periods of the vortex orbit, in contrast to the trajectories presented in [227, 140], in which the vortex generally precesses at a greater radius on each passage about the trap centre. In figure 7.2(b) we plot the vortex precession frequency. In practice, to evaluate this frequency, we follow the procedure of section 6.3.2, forming the covariance matrix (classical one-body density matrix)

$$G(\mathbf{x}, \mathbf{x}') = \langle \psi^*(\mathbf{x}) \psi(\mathbf{x}') \rangle_{\Omega}, \quad (7.2)$$

where $\langle \cdots \rangle_{\Omega}$ denotes a time average, in a frame rotating at angular velocity Ω about the trap axis. We vary Ω such that the largest eigenvalue of $G(\mathbf{x}, \mathbf{x}')$ is maximised. The value of Ω at which the maximum occurs is thus that of the rotating frame in which the coherent fraction of the classical field is most stationary, which forms a best estimate for the condensate angular frequency (vortex precession frequency) Ω_c , as discussed in section 6.3.2. As we consider a nonequilibrium scenario, we construct equation (7.2) by averaging over short time periods [29], in each case calculating the average of 250 consecutive classical-field samples over a ten cycle period. Figure 7.2(b) shows that the vortex precession frequency determined in this manner increases as the vortex displacement increases, as is well known for a zero-temperature condensate [87, 150] and observed in the simulations of [227, 140]. Moreover, we observe that the oscillations in vortex radius r_v at late times $t \gtrsim 2000$ cyc are accompanied by oscillations in the condensate angular velocity. The oscillations in the two are positively correlated, i.e., the decreases in vortex radius are associated with periods of slowing of the condensate rotation during its otherwise steady increase with time.

7.2.2 Rotational dynamics of condensate and cloud

The above definition of the condensate mode in terms of short-time covariance-matrix eigenvectors allows us to resolve the dynamics of the condensed and noncondensed com-




Figure 7.2: Quantities characterising the vortex motion. (a) Vortex radial displacement and (b) angular velocity of the condensate orbital. Quantities are for a simulation with trap anisotropy $\epsilon = 0.025$, corresponding to the classical-field densities in figure 7.1.

ponents of the field. As in section 6.4, we introduce the decomposition of the classical-field one-body density matrix in terms of its eigenvectors $|\chi_i\rangle$ and corresponding eigenvalues n_i

$$G = n_0|\chi_0\rangle\langle\chi_0| + \sum_{k>0} n_k|\chi_k\rangle\langle\chi_k| \equiv G_0 + G_{\text{th}}, \quad (7.3)$$

which separates it into condensed and noncondensed parts. Following section 6.4, we then define the averages of a single-body operator J in the condensed and noncondensed components of the field by $\langle J \rangle_0 = \text{Tr}\{G_0 J\}$ and $\langle J \rangle_{\text{th}} = \text{Tr}\{G_{\text{th}} J\}$ respectively. We use this decomposition to calculate the angular momentum of the condensate ($\langle L_z \rangle_0$) and thermal cloud ($\langle L_z \rangle_{\text{th}}$). These quantities are presented in figure 7.3(a), along with the total angular momentum of the field ($\langle L_z \rangle_{\text{tot}}$). We observe that in the first few hundred trap cycles after the introduction of the trap anisotropy, the angular momentum of the thermal cloud (solid line in figure 7.3(a)) undergoes approximately exponential decay, reaching values close to zero by $t \approx 1000$ trap cycles. By contrast, the angular momentum of the condensate (dashed line) is essentially unchanged during this period, which we identify with the initial period of quiescence of the vortex noted in section 7.2.1. At $t \approx 1100$ cyc, the angular momentum of the condensate begins to decay, and after $t \approx 3200$ cyc the condensate is essentially irrotational, with some small value of angular momentum remaining in the form of surface excitations. We note that the angular momentum of the thermal field component *increases* over the period $t \in [1000, 3200]$ cyc. We understand this as follows: initially the angular momentum of the thermal cloud is lost due to its interaction with the trap anisotropy. This makes the vortex thermodynamically unstable, and it thus begins to decay, liberating its angular momentum to the cloud. The cloud loses angular momentum to the trap at a rate proportional to its angular velocity (frictional loss $d\langle L_z \rangle/dt \propto -\Omega_{\text{th}}$ [276]), which is small immediately after the initial arrest of the cloud by the trap anisotropy. Due to its initially slow rotation, the cloud therefore gains angular momentum from the decaying vortex faster than it can dissipate it to the trap. As the thermal cloud's rotation rate increases, so does the rate at which it dissipates angular momentum to the trap, which acts to curb the increase in cloud angular velocity. However, the precessing-vortex condensate is thermodynamically anomalous: as it loses angular momentum, its angular velocity *increases* [42, 191, 150]. In the present case, the increasing rotation rate of the decaying vortex causes the rate of transfer of angular momentum from the vortex to the cloud (which is proportional to the differential angular velocity between the two [81]) to dominate the rate at which the cloud dissipates angular momentum to the trap, throughout the decay of the vortex. As shown in figure 7.3(b), this leads to ‘run-away’ spin-up of the vortex precession frequency, with the thermal-cloud rotation rate being driven up in sympathy.




Figure 7.3: Quantities characterising the temporal evolution of the condensate and thermal cloud. (a) Angular momenta of the condensate and thermal cloud, and total angular momentum of the field. (b) Angular velocities of the condensate and thermal cloud.

Closer inspection of the angular momenta $\langle L_z \rangle_0$ and $\langle L_z \rangle_{\text{th}}$ reveals negatively correlated oscillations in the angular momenta of the two components during the decay, as a result of the detailed nonlinear dynamics of angular momentum exchange and loss during the decay process. We identify this oscillation in angular-momentum transfer with the oscillations in vortex rotation rate discussed in section 7.2.1. In figure 7.3(b) we plot the angular velocity of the condensate Ω_c (dots) derived in the estimation of the condensate fraction (section 7.2.1), and the angular velocity of the thermal cloud $\Omega_{\text{th}} = \langle L_z \rangle_{\text{th}} / \langle \Theta_c \rangle_{\text{th}}$ (line), where we assume the expectation value of the classical moment of inertia $\Theta_c \equiv mr^2$ as an estimate of the cloud's true moment of inertia. In section 6.5.3 we discussed the level of uncertainty inherent in this procedure, nevertheless it yields a clear qualitative description of the decay dynamics. The oscillations in rotation rate of the thermal cloud are clearly visible here, and we note that they are *positively* correlated with the oscillations in angular velocity of the condensate, due to its anomalous rotational response. We note that oscillatory behaviour arises already in a linear analysis of the arrest of a rotating Boltzmann gas [120] by a trap anisotropy. It is therefore not surprising that similar oscillations occur in the transfer of angular momentum from the condensate to the nonequilibrium thermal field, considering the complexity of their coupled dynamics (cf. [276]), and the anomalous response of the vortical condensate.

Finally, we note that the cloud angular momentum (rotation rate) reaches its peak when the vortex leaves the condensate ($t \approx 3200$ cyc), after which it undergoes a second near-exponential decay phase, the angular momentum of the field decaying such that by $t \approx 5000$ cyc, only the thermodynamic fluctuations in $\langle L_z \rangle_{\text{th}}$ exhibited by the finite-temperature field at rest in the anisotropic potential remain.

7.2.3 Heating of the atomic field

We now consider the heating of the atomic field during the arrest of its rotation. As the system we evolve is Hamiltonian, with a time-independent potential, the total energy of the classical field is a constant of the motion. Consequently, the trap anisotropy dissipates the angular momentum of the field by converting the rotational kinetic energy of the field into internal energy [276], and we therefore expect some heating of the field to occur, due to this redistribution of energy. We can estimate the heating of the field as follows: the rotational energy of the gas is initially $(E_{\text{rot}})_i = \Omega_i L_i \approx 2.7 \times 10^3 \hbar \omega_r$. We compare this to the initial thermal energy, which should be reasonably well estimated by the energy added to the ground vortex state in forming the initial thermal state (section 7.1), $E_{\text{th}} \sim 0.1 E_g \approx 7.6 \times 10^3 \hbar \omega_r$. Assuming a linear relationship between thermal energy and temperature (which should be valid for the low temperatures we consider here, see [64]),

we might therefore expect an increase in the field temperature of $\sim 30\%$ during the arrest of the condensate rotation.

In order to quantify the heating and its development during the field evolution, we estimate the (effective) thermodynamic parameters of the field (chemical potential μ and temperature T) using the fitting procedure of appendix D. As in section 6.5.4, we allow for the differential rotation between the thermal cloud and the frame in which the classical-field cutoff is effected (the laboratory frame, in the present case). We perform a fit to the wing of the (time-averaged) radial distribution $n(r)$ of the classical field; i.e., we fit over the radii range $r \in [r_- + 0.5r_0, r_{\text{tp}}]$ where r_- marks the minimum of the Hartree-Fock effective potential $V_{\text{eff}} = (m/2)[\omega_r^2 - \Omega_{\text{th}}^2]r^2 + 2U_{2\text{D}}n(r)$ ¹, and r_{tp} is the semiclassical turning-point of the low-energy space \mathbf{C} . We fit to radial densities of the classical field averaged over 10 cyc periods, and assume the values of Ω_{th} calculated in section 7.2.2. The resulting estimates for μ and T are presented in figure 7.4(a). We find that the temperature (solid line) is approximately constant during the first ~ 1000 cycles of the field evolution, corresponding to the initial arrest of the thermal field component. The initial angular momentum of the thermal component is small (see figure 7.3(a)), comprising some $\sim 25\%$ of the total angular momentum of the field, and so any heating of the field due to the redistribution of the associated rotational energy is possibly too small to resolve above the uncertainty in the temperature estimates. Beginning at $t \approx 1100$ cyc, corresponding to the start of the vortex-decay phase, the temperature exhibits a steady, approximately linear increase. We associate this increase with the conversion of rotational kinetic energy to thermal energy, due to the action of the trap anisotropy, and presumably also as a result of the scattering of excitations by the vortex which produces the frictional effect [81]. There is a final sharp rise in temperature at $t \approx 3000\text{--}3200$ trap cycles. This rise is perhaps due to the final decay of the vortex into excitations at the surface of the condensate [81]. We note however that the field, and in particular the condensate surface, undergoes strong fluctuations as the system passes through this transition to the vortex-free state, with, e.g., multiple (ghost) vortices present at the condensate surface at times. This can be viewed as the system essentially ‘reversing’ through the surface-wave instability arising from the relative motion of condensate and thermal cloud, in which vortices spontaneously grow from surface excitations [11, 266, 200]. The thermodynamic parameters may therefore be ill-defined during this period. After this period of strong surface fluctuations, the temperature levels off at $T \approx 21\hbar\omega_r/k_{\text{B}}$, corresponding to heating of $\approx 50\%$ during the arrest. We note that the chemical potential (dashed line in figure 7.4(a)) exhibits a slight downward trend beginning around $t \approx 1100$ cyc, falling by $\sim 0.4\hbar\omega_r$, though we note that

¹We refer here to the local minimum associated with the boundary of the condensate, rather than any minimum that may occur due to the presence of the vortex core.

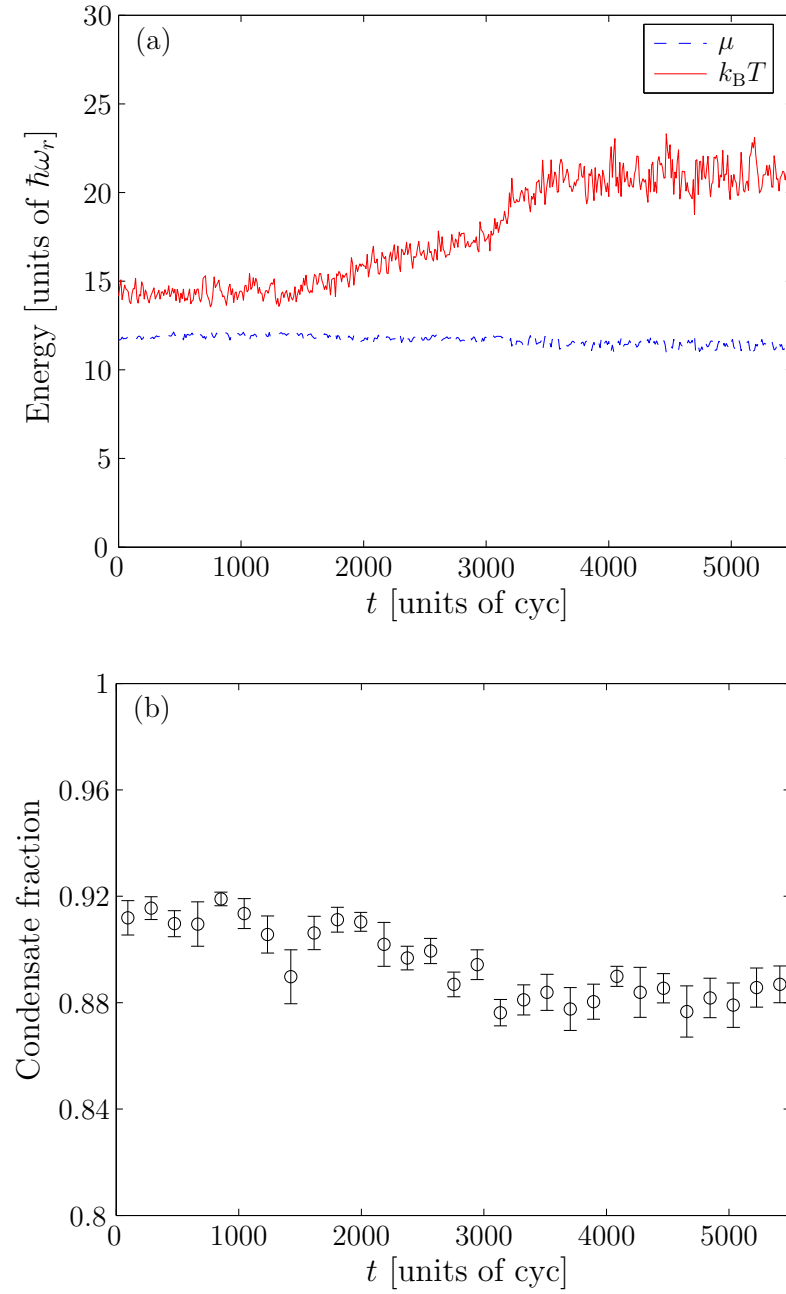


Figure 7.4: Heating of the field. Evolution of (a) the effective chemical potential and temperature and (b) the condensate fraction.

this change is of the same order as the variation in estimates for μ obtained at late times $t \gtrsim 3500$ cyc.

In figure 7.4(b) we plot the condensate fraction $f_c \equiv n_0 / \sum_k n_k$ obtained from the procedure outlined in section 7.2.1. The fluctuations in estimates of this quantity are large, as is expected given the short time scale (10 cyc) over which the appropriate averages are taken, and the nonequilibrium nature of the field. We therefore group the condensate-fraction estimates into bins each containing 19 consecutive estimates of f_c , and average over the estimates in each bin (cf. section 6.3.3). We plot the mean and standard deviation of the estimates in each bin (circles and their error bars), and observe a clear downwards trend as time proceeds, with the condensate fraction dropping by $\sim 3\%$ during the arrest.

7.3 Dependence on trap anisotropy

We now consider the effect of varying the trap anisotropy ϵ on the behaviour of the classical-field trajectories. We intuitively expect the rate at which the trap dissipates the angular momentum of the thermal cloud to depend strongly on the magnitude of the anisotropy, and a quantitative model for this dependence in the case of a classical (Boltzmann) gas was presented in [120]. By contrast, the efficiency with which the thermal cloud extracts angular momentum from the condensate is dictated by the (longitudinal) mutual-friction coefficient, which depends on the temperature and chemical potential of the field, in addition to its microscopic properties [125, 244, 81]. We therefore expect, as discussed in the vortex-continuum analysis of [276], to explore different regimes of relaxation dynamics as we vary the trap anisotropy and consequently the relative strengths of vortex-cloud and cloud-trap friction. In this section we consider the evolution of the classical field in the presence of anisotropies of different magnitudes, and observe the resulting differences in the behaviour of the field.

7.3.1 Weak damping

In figure 7.5(a) we plot the evolution of the angular momentum of the condensate and thermal cloud in the presence of a weak trap anisotropy ($\epsilon = 0.005$). We observe that after an initial decline over the first $\approx 10^4$ trap cycles, the cloud angular momentum, despite exhibiting large fluctuations, maintains a reasonably steady mean value which increases slowly over the period $t \sim [10^4, 3 \times 10^4]$ trap cycles. During this period the angular momentum of the condensate drops by approximately a factor of 2, and large fluctuations in the distribution of angular momentum between the two components are visible. This behaviour is again apparent in the vortex precession frequency and cloud




Figure 7.5: (a) Angular momenta and (b) angular velocities of field components in a simulation with trap anisotropy $\epsilon = 0.005$. (c) Angular momenta and (d) angular velocities of field components in a simulation with trap anisotropy $\epsilon = 0.100$. Inset: Evolution of the condensate fraction in the case $\epsilon = 0.100$.

rotation rate presented in figure 7.5(b). This shows that the cloud exhibits a steady-state nonequilibrium behaviour, i.e., its rotation rate remains somewhere between that of the precessing condensate mode, and that of the trap ($\Omega_{\text{tr}} = 0$). In this regime its angular momentum remains approximately constant, as the rate at which it gains angular momentum from the decaying vortex matches the rate at which it loses angular momentum to the trap. We expect the angular velocity at which this balance occurs, and indeed whether such a regime occurs at all, to depend strongly on the relative strengths of the vortex-cloud and cloud-trap friction. We note that the angular velocity of the cloud appears to slowly increase over time, as the condensate loses angular momentum and its angular velocity increases, shifting the cloud rotation rate at which the angular momentum transfer rates are balanced. At $t \approx 3.2 \times 10^4$ cyc the vortex's precession accelerates as it approaches the condensate boundary, and during the period $t \in [3.2 \times 10^4, 3.4 \times 10^4]$ cyc the adiabaticity of the vortex dissipation is lost as the field enters the critical regime associated with nucleation of the vortex at the surface (see section 7.2.3). In this regime the condensate surface is unstable and as noted in section 7.2.3, multiple 'ghost' vortices may be present simultaneously at the condensate boundary. A large amount of angular momentum is transferred nonadiabatically to the thermal cloud during this period of criticality, while small surface-mode oscillations of the condensate persist until $t \approx 4 \times 10^4$ cyc. The angular momentum is subsequently dissipated from the thermal cloud over a period of $\sim 3 \times 10^4$ cyc, in an oscillatory but approximately exponential decay phase.

7.3.2 Strong damping

We now turn our attention to a scenario of vortex arrest due to the presence of a strong trap anisotropy ($\epsilon = 0.1$). The dynamics in this case are strongly nonequilibrium and the arrest of the field's rotation occurs on a shorter time scale than the cases already considered, and so here we form the density matrix (7.2) by averaging classical-field samples over a shorter period of two trap cycles. In this case the angular momentum of the thermal cloud (solid line in figure 7.5(c)) is rapidly depleted, and actually fluctuates below zero by $t \approx 400$ trap cycles. The condensate is slower to respond; during this period the condensate angular momentum decreases by $\sim 7\%$. Subsequently, the cloud angular momentum fluctuates strongly, and at $t \approx 750$ cyc rises to $\sim 2000\hbar$ (close to its initial value), as the vortex is rapidly expelled from the condensate. The angular momentum is then dissipated from the cloud over a period $\lesssim 100$ cyc, and thereafter the angular momentum of the field fluctuates about zero. The rapid expulsion of the vortex, and strong fluctuations of the cloud rotation, are again visible in the calculated angular velocities (figure 7.5(d)).

The behaviour of the cloud angular momentum in this trajectory suggests that the cloud is overcritically damped by the trap anisotropy [120]. Its response to the anisotropy appears unhindered by its coupling to the condensate via the vortex core, and the angular momentum it acquires from the vortex is quickly yielded to the trap, despite its small angular velocity $\Omega_{\text{th}} \lesssim 0.1$. The relaxation process in this case is violently nonadiabatic, as evidenced by the strong fluctuations in the cloud angular momentum. Indeed during the period $t \sim [600, 700]$ cyc large, long-wavelength surface oscillations are visible in the field as the vortex precesses rapidly near the condensate boundary. It appears that the strong trap anisotropy and rapid ‘stirring’ motion of the vortex conspire to strongly perturb the condensate bulk in this regime. During this period the measured condensate fraction is suppressed (inset to figure 7.5(d)) due to the strongly nonequilibrium behaviour of the condensate, in which surface-wave excitations define frames of rotation distinct from that of the vortex. The decomposition of the field into condensed and noncondensed components must therefore be viewed with some caution in this strongly nonequilibrium scenario. Nevertheless, it is clear that the effect of the trap on the cloud dominates the vortex-cloud coupling in this scenario, in stark contrast to the near-adiabatic decay scenario of section 7.3.1.

7.3.3 Decay times

We now consider the dependence of the relaxation times on the strength of the trap anisotropy. As the relaxation of both the condensate and the thermal cloud is generally nonexponential, we consider the times at which the condensate becomes irrotational, and at which the total angular momentum of the field is lost. Due to the persistence of surface excitations which prevent the condensate angular momentum from reaching zero, we define the condensate stopping time τ_{cond} as that at which the angular momentum of the condensate first drops below $0.02\hbar$ per particle, and the field stopping time as the first *subsequent* time τ_{field} at which the cloud angular momentum crosses zero. These times are presented in figure 7.6. We observe that the condensate arrest times (circles in figure 7.6) vary by two orders of magnitude over the range of trap anisotropies we consider. Moreover, the total-field relaxation times (plusses) become increasingly longer than the vortex relaxation times as the trap anisotropy is weakened, causing slower dissipation of the cloud angular momentum even in the absence of the vortex. We therefore consider the time $\Delta\tau = \tau_{\text{field}} - \tau_{\text{cond}}$ over which the angular momentum of the cloud dissipates following the expulsion of the vortex. Although the angular momentum lost in this final damping phase varies over the range $\langle L_z \rangle_{\text{th}} \sim 3000\text{--}5000\hbar$, precluding a precise analysis, the times $\Delta\tau$ provide a useful characterisation of the dependence of the cloud relaxation




Figure 7.6: Relaxation of the field. Circles (plusses) represent the times at which the vortex leaves the condensate (total field angular momentum reaches zero). Inset: Cloud spin-down times and the linear fit performed to extract their scaling with trap ellipticity.

on the trap anisotropy. We perform a linear fit to $\Delta\tau$ as a function of ϵ in log-log space (inset to figure 7.6), and find the scaling $\Delta\tau \propto \epsilon^{-2.1}$, in good agreement with the scaling $t_{\text{down}} \propto \epsilon^{-2}$ for thermal-cloud spin-down predicted for weak anisotropies by Guéry-Odelin [120]. The latter calculations considered a simpler scenario in which no condensate was present, however once the condensate becomes vortex-free its primary influence on the thermal cloud is simply to reshape the trapping potential experienced by the thermal atoms. The quantitative agreement between the spin-down scaling laws predicted by these two very different approaches to describing the thermal-cloud dynamics is a strong indication of the utility of the (superficially Gross-Pitaevskii-like) classical-field model for modelling the dynamics of *thermal* atoms in the degenerate regime.

7.4 Summary

We have carried out the first simulations of the arrest of a rotating Bose-Einstein condensate due to the presence of a trap anisotropy which includes the coupled nonequilibrium dynamics of the condensate and thermal cloud. Our method makes no assumptions of stationarity of a thermal bath, nor are the two components artificially given disparate

rotational parameters. Rather, our approach describes the dynamical migration of an *equilibrium* rotating thermal state to a new, irrotational equilibrium, solely due to the action of the trapping potential.

We observe for all parameters we considered that the rotation rates of the condensate and thermal component are distinct during the decay. The anomalous rotational response of the precessing-vortex condensate can lead to a counter-intuitive spin-up of the thermal component during the decay, and we observe nonequilibrium oscillations in transfer of angular momentum between the two components as the vortex responds to the dissipative effect of the thermal cloud, which is itself damped by the trapping potential anisotropy.

For trap anisotropies that are weak, the thermal field settles to a nonequilibrium steady state, with rotation rate intermediate between that of the condensate and that of the (static) trap. In this scenario the angular momentum of the condensate is slowly depleted while that of the thermal cloud remains nearly constant, until the vortex nears the condensate boundary and the linearity of the vortex decay breaks down. For stronger trap anisotropies, the angular momentum of the thermal cloud may be almost entirely depleted before the condensate responds.

We quantified the heating of the atomic field during the arrest, and found it to be commensurate with the conversion of rotational kinetic energy into thermal energy by the trap anisotropy. We also considered the time scales over which the condensate angular momentum and total field angular momentum were dissipated, and found reasonable quantitative agreement between the scaling of the thermal cloud spin-down time with trap anisotropy and the predictions of a Boltzmann-gas model [120].

Chapter 8

Prior studies of vortex-lattice formation in stirred condensates

In this chapter we review the key prior work on the stirring of Bose-Einstein condensates to produce vortex lattices. As noted in section 1.2, a range of early experiments considered the formation of individual vortices, by techniques including coherent hyperfine-component interconversion [180, 7], the ‘sweeping’ of a laser light-shift potential through the atom cloud [136], and phase imprinting [159]. Other experiments considered the formation of vortex lattices by setting a normal thermal cloud into rotation in cylindrically symmetric trap, and cooling it through the condensation transition [124]. However, we here consider the formation of vortices in condensate *stirring* experiments, in which an anisotropic rotating trapping potential defines a rotating frame of equilibrium for the atomic cloud, and induces the formation of vortices in the condensate. This approach was used to produce the first multiple-vortex configurations and vortex lattices in condensed Bose gases, and as we discuss in this chapter, exhibits a rich phenomenology. Indeed, the experimental investigation of vortex nucleation and lattice formation by stirring has stimulated substantial theoretical research over recent years, and led to some controversy over the precise mechanisms of the vortex nucleation process, and the role of thermality in the formation of a rotating vortex lattice. This chapter is structured as follows: in section 8.1 we review the results of experiments on condensate stirring performed by the groups of Dalibard (École Normale Supérieure), Ketterle (MIT), and Foot (Oxford). In section 8.2 we review the various theoretical approaches to explaining the experimental results, and in section 8.3 we formulate the issues on the nature of the vortex formation which remained unanswered by those works. We will address those questions with our own classical-field methodology in chapter 9.

8.1 Experiments

8.1.1 Work of Dalibard's group (ENS)

The group of Dalibard performed the first experiments aimed at reproducing the ‘rotating-bucket’ paradigm for studies of rotation in superfluid helium. In the helium experiments (see [75] and references therein), the rotating bucket defines the frame of thermal equilibrium through its surface roughness; as the rotation rate of the bucket increases, the superfluid admits vortices which form a regular array or *vortex lattice* at equilibrium. By analogy, a rotating frame can be imposed on a gaseous condensate by deforming the trapping potential away from cylindrical symmetry, and rotating it. In the experiments performed on ^{87}Rb atoms by Madison *et al.* [175], this deformation was provided by the light-shift potential provided by a blue-detuned laser beam (see, e.g., [44]), which was *toggled* between two diametrically opposite focus points in the transverse plane of the trap, at a frequency very rapid compared with the dynamics of the atomic mean field. The resulting effect was that the prolate trapping potential was (to leading order) anisotropically deformed in the xy plane, leading to an elliptical cross-section. This deformation of the trapping potential was then rotated about the trap axis at a frequency Ω . In the first experiments [175, 176], this stirring potential was switched on during the final stages of evaporative cooling, at the end of which, the rf shield was left turned on. The atomic cloud was left to thermalise in the rotating ‘bucket’, and the resulting structure of the condensate was then studied by time-of-flight analysis. The somewhat surprising result of these experiments was the high stirring frequency at which vortices were first observed, which was significantly higher than the *thermodynamic* critical frequency Ω_c at which the energy of the vortex state becomes lower than that of the vortex-free state (section 2.3.2). Dalibard's group then undertook further experiments [52], in which the stirring was performed on a preformed condensate *after* the completion of the evaporative cooling, and they considered in detail the dependence of the vortex nucleation on the ellipticity of the trapping potential, which they varied by adjust the intensity of the stirring laser beam. They found that for small ellipticities, vortices were formed in a resonant regime centred on the resonant frequency $\omega_Q = \omega_r/\sqrt{2}$ of the quadrupole surface mode (see section 2.2.2). Furthermore the frequency width of this regime narrowed with decreasing ellipticity ϵ , indicating that the resonant excitation of the quadrupole mode was responsible for the processes leading to vortex nucleation. Subsequent investigations performed by these authors [174] into the *vortex-free* rotating ellipsoidal configurations of the condensate predicted by Recati *et al.* [220], gave further support to this conclusion.

8.1.2 Work of Ketterle’s group (MIT)

Ketterle’s group at MIT performed similar experiments to those of the ENS group, stirring ^{23}Na condensates and creating large lattices (with $\gtrsim 100$ vortices) [2]. They used a two-laser pattern to stir the condensate, similarly to the ENS experiments, however the resulting optical potential corresponded to a strongly anharmonic deformation of their trap potential. In these experiments, the power of the stirring laser was ramped on, held for a variable hold time, and then ramped off. They observed that for shorter stirring times ($t_{\text{stir}} \lesssim 500\text{ms}$) the number of vortices nucleated was resonantly enhanced near the frequency ω_Q of the quadrupole mode resonance, as observed in the ENS experiments. They also observed a broadening of this resonance (with respect to the stirring frequency) with increasing stirring time t_{stir} , and noted an increase in the number of vortices formed *far* from this resonance as a function of the stirring time. Vortices were observed to form at drive frequencies as low as $\approx 0.27\omega_r$, which was however large compared to the thermodynamic stability frequency $\Omega_c \approx 0.08\omega_r$ for their system.

In the formation of large lattices, these authors noted the emergence of a blurry nonequilibrium structure at short times after the stirring, with density features they attributed to misaligned, disordered vortices. Taking time-of-flight images after increasing wait times, they observed the increasing alignment and ordering of the vortices as time progressed. They noted also that off-resonance stirring of the condensate resulted in lattices with fewer vortices, which took longer to equilibrate into a crystalline structure. These authors argued that the minimum frequencies for vortex nucleation they observed were in fair agreement with the Landau instability frequency (section 2.3.1) of their condensate (corresponding to surface modes with angular-momentum projection $m \approx 18$, yielding $\Omega_L \approx 0.3\omega_r$). They noted however that their anharmonic stirring potential leads to strong deformations of the condensate over a wide range of frequencies, and thus allows for the excitation of large- m surface modes, which significantly complicates the analysis of the nucleation process in these experiments.

8.1.3 Work of Foot’s group (Oxford)

Foot’s group studied the nucleation of vortices by stirring a condensate of ^{87}Rb atoms with a purely magnetic trapping arrangement [130]. They achieved the (approximately) elliptic deformation of their oblate time-orbiting-potential trap by modulating the x and y components of the applied magnetic bias fields. In practice, they first condensed the Bose cloud in a ‘static’ (cylindrically symmetric) trap not subject to any rotation, achieving temperatures of $\sim 0.5T_c$. They then ramped on the *ellipticity* of the constantly rotating trap deformation from zero to some final value over a period of time (200 ms). In contrast

to the ENS and MIT experiments, the evaporative rf shield was *not* engaged during the stirring process. They reported significant heating of the atomic field during the stirring process, with the final temperature rising to approximately $0.8T_c$ in the absence of the rf shield. For a given trap ellipticity, they observed that the nucleation of vortices occurred for a narrow range of rotation frequencies centred approximately around the quadrupole resonance frequency ω_Q . In agreement with the ENS results, they found the breadth of the resonance increased with increasing ellipticity ϵ . In further investigations, they fixed the drive frequency, and ramped the trap anisotropy on adiabatically, and thus measured the critical anisotropy for vortex nucleation at fixed angular velocity. In contrast to the results of Madison *et al.* [175, 52], they observed the nucleation of vortices for drive frequencies $\Omega < \omega_Q$, with an increasing ellipticity required for vortex nucleation as the drive frequency was reduced further below the resonance frequency. They suggested several possible explanations for this phenomenon: that (i) it was due to the thin tails of the instability locus predicted by the theory of Sinha and Castin [240] (discussed in section 8.2), or; (ii) higher-order anisotropies or anharmonicities of the trap were present, leading to the excitation of other surface modes (cf. [52]), or; (iii) the presence of the thermal cloud was responsible for facilitating the Landau instability in this regime. They investigated the effect of temperature by conducting experiments with the rf shield left on, mitigating the heating of the condensate during the stirring process, obtaining a final temperature $T \approx 0.5T_c$, and observed no significant change in the nucleation curve. They did note, however, that the presence of the shield had a significant impact on the relaxation and decay of the vortex lattices, with stable lattices being much more likely to form and persisting for longer in the presence of the rf shield.

8.2 Theory

Work of Feder *et al.*: Anomalous modes, and GP simulations of stirring

Feder *et al.* [80] identified that while global thermodynamic considerations show that a vortex-containing state is the ground state for rotation rates exceeding some frequency Ω_c , the imposition of a rotation at this frequency is not necessarily sufficient to lead to vortex nucleation. They noted that in general, there is a higher frequency Ω_v at which surface excitations become unstable to growth by the Landau instability (section 2.3.1). Furthermore, in the ENS experiments, the frequency at which vortices are first observed is higher again than Ω_v . In their paper [80] they focussed on a third relevant frequency: the frequency Ω_m at which a condensate containing a vortex first becomes *metastable* (section 2.3.3). In the Bogoliubov approach to the description of the excitations of the

vortex, this is the rotation frequency at which the anomalous mode(s) of the vortex state (which have negative energies in the lab frame *by definition*) are all Doppler shifted to positive values by the rotation. They noted that for prolate traps (such as those of the ENS experiments), additional anomalous modes (associated with deformations of the vortex line) appear, and can lead to a metastability frequency Ω_m which can exceed both Ω_v and Ω_c . They conjectured that the comparatively high rotation rate required in the ENS experiments was due to the instability of the vortex state below this metastability frequency. However, subsequent studies of vortex formation in highly prolate condensates by Dalibard's group [174] found no evidence of the inhibition of vortex nucleation due to anomalous modes.

Feder and his colleagues also performed the first *time-dependent* Gross-Pitaevskii equation simulations of the stirring process, evolving a condensate corresponding to the ENS parameters in the presence of an abruptly introduced, rotating trap anisotropy. They performed two such simulations (performed in the frame of the rotating trap), with drive frequencies $\Omega = 0.7\omega_r$ and $\Omega = 0.8\omega_r$ respectively. In the former case, they found no nucleation of vortices, while in the case $\Omega = 0.8\omega_r$, they observed vortices nucleating at the condensate surface and migrating into the centre of the atomic cloud. However, the motion of the vortices remained turbulent, and the authors concluded that dissipative processes beyond the GPE may be required to facilitate the crystallisation of the vortex array to a rigid lattice.

Work of Sinha and Castin : Dynamical instabilities of the condensate

Sinha and Castin made a significant advance [240] in the understanding of the role of quadrupole oscillations in the stirring experiments, identifying the *dynamical instability* of this surface mode as crucial in the vortex nucleation. They considered a classical hydrodynamic description [250] of the condensate flow pattern in the regime of quadrupole oscillations, from which they determined numerically the conditions (rotation frequency and trap anisotropy) for which the quadrupole oscillations became dynamically unstable (see section 2.2.3). They considered two distinct stirring scenarios: (i) an adiabatic scenario, in which the ellipticity of the trap is held constant, while the rotation frequency Ω of the anisotropy is ramped up adiabatically from zero; (ii) abruptly turning on the stirrer at a constant rotation frequency. In the first scenario, the condensate adiabatically follows a steady-state rotating-ellipsoid pattern, as previously predicted by Recati *et al.* [220]. The flow pattern of the equilibrium condensate in this regime is described (in the hydrodynamic limit) by quadratic ansätze for the density and phase, and so these authors characterised the dynamic stability of the condensate mode in the rotating trap

by considering the (linearised) evolution of small deviations about this flow pattern, employing polynomial ansätze for the density and phase deviations of order $n = 3, 4, \dots$. Constructing and diagonalising the evolution operator for these deviations, they identified the parameters for which modes with eigenvalues λ with positive real parts, corresponding to dynamically unstable excitations of the condensate, appeared. They found that the locus of points in the Ω - ϵ plane for which the condensate is unstable forms a pattern comprised of several crescent structures¹ with, in particular, a broad base in the limit $\epsilon \rightarrow 0^+$, occurring for $\Omega > \omega_Q$. In the second scenario, which corresponds to the (later) experiments of the ENS group [52], the ellipticity of the stirrer (with fixed rotation frequency) is abruptly turned on. Adiabatic following is not applicable in this case, so they calculated the time-dependent evolution operator for the linearised deviations as a function of time and considered the growth of the largest eigenvalue of this operator with increasing time t . They found that in this abrupt scenario, the quadrupole oscillations were dynamically unstable in a narrow range of Ω , given approximately by $\Omega \in [0.71, 0.74]$ for an ellipticity $\epsilon = 0.01$. These results agree well with those of ENS experiments [52], strongly supporting the hypothesis that the dynamical instabilities of the quadrupole oscillations are instrumental in the formation of vortices.

Work of Tsubota *et al.*: The role of dissipation

Tsubota *et al.* [254] noted the necessity of *dissipation* to allow the condensate to relax to the vortical state, which is of a lower energy (in the frame co-rotating with the trap) than the initial vortex-free state. They therefore considered the solution of a modified GPE, which includes a *phenomenological* damping term, following the earlier work of Choi *et al.* [54]. In essence, this approach amounts to evolving the time-dependent GP in *complex time*, which hybridises the standard real-time evolution of the equation with the imaginary-time approach to finding ground states of the time-independent GP equation (see, e.g., [61]). In general, such a damping term leads to a change in normalisation of the wave function, however these authors continuously adjusted the chemical potential appearing in the GP equation throughout the evolution, in order to ensure conservation of the GP-wave-function normalisation. These authors used this modified GP equation to simulate the stirring of the condensate with an elliptical potential and observed a transition to the vortex-lattice state. However, they found different driving frequencies for the onset of lattice formation depending on the stirring procedure, namely $\Omega = 0.57\omega_r$ for the ‘abrupt’ driving scenario and $\Omega = 0.75\omega_r$ in the ‘adiabatic’ scenario. While

¹Anglin [12] has shown that the appearance of such unstable eigenvalues in narrow ‘finger-like’ regions of parameter space is a generic feature of dynamically unstable interacting many-body systems.

it seems clear that the latter case involves the dynamical instability of the quadrupole mode, these authors offered no interpretation for the lower frequency they found in the ‘abrupt’ scenario.

Work of Williams *et al.*: ZNG theory for surface-mode growth

Subsequently to the work of Tsubota *et al.* [254], Williams *et al.* [266] considered the dissipative effects of coupling of the condensate mode to the thermal cloud within the ‘ZNG’ Gross-Pitaevskii-Boltzmann theory [274]². They derived analytical results for the damping of a surface mode of angular-momentum projection m along the trap axis, due to the C_{12} term of the ZNG theory, which describes collisions between condensed and noncondensed atoms which result in the transfer of atoms into (or out of) the condensate. They found that in the presence of a rotating thermal cloud, the damping of surface modes which ‘rotate’ with the opposite sense to the cloud’s circulation increases linearly with the rotation frequency of the cloud, while the damping of modes that rotate with the *same* sense as the cloud *decreases* with increasing Ω . Moreover, they found that the damping of a given mode with angular momentum projection $m > 0$ changes sign (so as to become *growth*) for a cloud driving frequency $\Omega_m = \omega_m/m$, i.e., they derive explicitly the result that the onset of Landau instability (negative excitation frequency) of a surface mode in the GP formalism leads to the spontaneous *growth* of this mode in the presence of a thermal cloud rotating above the corresponding Landau frequency Ω_m ³.

Work of Penckwitt *et al.*: SGPE theory for lattice growth and stabilisation

Penckwitt *et al.* [200] considered the problem of vortex nucleation within the SGPE framework of Gardiner *et al.* [102]. Their model includes the dissipative effect of collisions between the condensate and noncondensed atoms, similarly to the ZNG model of Williams *et al.* [266], providing a more rigorous basis for a damped-GPE approach such as that of Tsubota *et al.* [254]. Their thermal atoms are treated as a bath of fixed temperature, chemical potential, and angular velocity, and the latter can be distinct from that of the rotating trapping potential. In contrast to the work of Tsubota *et al.*, the chemical potential appearing in the evolution equation used in this work has the physical meaning of the chemical potential of the thermal cloud, and the condensate population grows (or shrinks) until the condensate chemical potential matches that of the bath. Penckwitt *et al.* focussed primarily on the scenario of a vortex lattice forming as a rotating thermal cloud

²We note that this formalism is reviewed (along with many others) in great detail in the tutorial review of Proukakis and Jackson [217].

³We note, however, that their calculation also reveals a shift of the relevant mode frequencies ω_m from their $T = 0$ values due to the coupling to the thermal cloud.

condenses in a cylindrically symmetric trap, which was demonstrated experimentally by Haljan *et al.* [124]. In addition to their numerical simulations of this growth scenario, they found analytically from their model that growth occurs for surface modes satisfying the Landau instability criterion, and furthermore derived an analytic result for the gain profile of the surface modes, finding good agreement with their numerical results. Turning their attention to the condensate stirring, they considered the ‘abrupt’ stirring scenario, with a drive frequency $\Omega = 0.65\omega_r$, both in the absence and the presence of a thermal cloud rotating at the same angular velocity. They found coherent, stable quadrupole oscillations of the condensate in the former case, and the nucleation of vortices in the presence of the cloud. On the basis of their results, the authors conjectured that the nucleation of vortices and their stabilisation into a rigid lattice is driven by the rotating thermal cloud, and that as such, the principal effect of the rotating trap anisotropy in experiments is to generate such a thermal cloud.

Work of Kasamatsu *et al.*: Separating the roles of the instabilities

Following on from their earlier work [254], the group of Ueda revisited the condensate-stirring scenario in a second, longer article [145]. In the first part of their paper, they considered again the quadrupole oscillations of the condensate in the presence of the rotating elliptical trap anisotropy. Adopting the quasiparticle-projection method of [184], they considered the evolution of the surface modes of the condensate in the presence of the rotating drive. The interpretation of the quasiparticle projection technique is somewhat dubious in this regime, as (i) the depletion of the initial circularly symmetric condensate mode is substantial, invalidating a linearised treatment and (ii) the coherent quadrupole oscillations described by the GPE correspond to the *collective* excitation of the condensate, i.e., a deformation of the condensate mode rather than the population of quasiparticle modes about it (see, e.g., [50]). Nevertheless, their procedure does yield a meaningful projection onto the relevant angular momentum components. They show that driving of the quadrupole mode at frequencies below its resonance leads to *recurrence* of the initial condensate state, while driving at higher frequencies leads to *stochastisation* of the GP evolution [237], in which case the trajectories become irregular and chaotic. They note that this is analogous to the well-known threshold behaviour of chaos and ergodicity in some nonlinear systems, as famously discovered in the pioneering numerical investigations of Fermi, Pasta, and Ulam [84]. They therefore conjectured that the role of the dynamical instability identified by Sinha and Castin [240] is to generate such stochastisation of the condensate evolution.

Armed with the interpretation of the phenomenological damping term furnished by the work of [266] and [200], they considered again the evolution of the condensate mode in the damped-GPE approach. They concluded from these investigations that the Landau instability of surface modes in the presence of the rotating thermal cloud is the fundamental mechanism of vortex nucleation (and, in particular, the origin of the critical frequency $\Omega = 0.57\omega_r$ observed in their earlier simulations [254]). They then considered the relation between the Landau instability of surface modes and the dynamical instability of the quadrupole oscillations of the condensate, and conjectured that the role of the dynamical instability is to generate a thermal component of the field, thus explaining the later results of the ENS group [52]. They noted that this also explains why vortex nucleation was occasionally observed at significantly lower frequencies: it is possible that vortex nucleation may occur simply due to Landau instability driven by the (small) thermal cloud already present in the experiments, but that this process is very slow, and in general the excitation of the dynamical instability is therefore necessary to form a sufficient thermal component to allow the condensate to relax to a vortical state. In their summation, they identified the need for a framework beyond the Gross-Pitaevskii theory in order to describe the dynamics of vortex nucleation as driven by a dynamically generated thermal component, arising from the dynamical instability of the quadrupole mode⁴.

Work of Lundh *et al.*: GP investigations of sub-resonant nucleation

Lundh *et al.* [171] performed systematic investigations of the vortex nucleation with time-dependent GPE calculations. They considered the scenario of [130] in which the ellipticity of the constantly rotating trap is ramped on adiabatically, and measured the critical ellipticity for vortex nucleation as a function of the drive frequency. At frequencies *above* the hydrodynamic resonance frequency ω_Q of the quadrupole surface mode, they found good agreement with the hydrodynamic predictions of references [220, 240]. They also found reasonable agreement with the results of [130] for the critical ellipticities at sub-resonant stirring frequencies ($\Omega < \omega_Q$). They suggested that vortex nucleation in this regime resulted from a mechanism similar to that discussed by Krämer *et al.* [153]: under quadrupole deformation of the condensate, the barrier to vortex nucleation (section 2.3.2) admits a saddle point which may become lower in energy than the isotropic vortex-free state, allowing the system to pass to the lower energy vortex state. However, we note that this effect only addresses the *thermodynamic* stability of the state; i.e., dissipation would still be required in order to allow the condensate to pass to the vortex state. It is

⁴Indeed, as the dynamical instability of the condensate involves (by definition) the exponential growth of noncondensate modes, the onset of the dynamical instability signifies the breakdown of the Gross-Pitaevskii theory (see section 2.2.3).

in our opinion more likely that the dynamical instability mechanism [240] underlies the sub-resonant vortex nucleation observed in their GP simulations and in the experiments of Hodby *et al.* [130].

These authors also consider the dynamics of the stirred field over longer time scales, and obtain results similar to those of Feder *et al.* [80], with a turbulent field developing after the instability, and the vortices failing to crystallise on the time scales they consider.

Work of Lobo *et al.* : Dissipation from turbulence

The work of Lobo *et al.* [170] was the first attempt, using a framework beyond the phenomenologically damped Gross-Pitaevskii theory, to describe the dynamics of vortex nucleation following the stirring of an (initially) zero-temperature condensate. Following from their earlier work on the truncated-Wigner formalism [238, 239], these authors considered the evolution of a classical field under the action of the Gross-Pitaevskii *equation*, but interpreted the field as representative of the whole atomic field (see chapter 3). In this approach, the turbulent nonlinear dynamics of the classical field provides for the *internal* dissipation of the energy liberated by the condensate during the vortex nucleation. These authors considered the evolution of a prolate condensate on a 3D Cartesian grid, in the frame of the trapping potential, whose rotation frequency is adiabatically ramped up. They observed the onset of the dynamical instability of the initial GP mode, and the subsequent evolution of the field into a turbulent state, after which vortex lines were nucleated into the condensate, and ultimately relaxed to a lattice configuration co-rotating with the drive. The authors estimated the temperature of the field at the end of their simulations, by assuming that the excess energy of the field over that of the ground lattice state was equilibrated over weakly interacting (Bogoliubov) modes. They found that the temperature was very small, of the order of one-tenth of the chemical potential of the field. This paper thus provided a significant advance in understanding of the problem: by considering the field ψ evolved by the GPE as the whole atomic field, rather than merely the condensate wave function, the dynamical instability of the condensate and the subsequent thermalisation dynamics of the atomic field could be studied within a unified framework. Moreover, their results suggest that the turbulent states obtained in the simulations of [80] and [171] are in fact *finite-temperature* states, beyond the validity of the GP approach. These authors also considered the stirring of initially finite-temperature classical-field configurations, and observed vortex nucleation for stirring above the thermodynamic critical frequency Ω_c .

Work of Parker and Adams: Turbulence and symmetry breaking

Subsequent work on the scenario of vortex nucleation by stirring with an elliptical trap anisotropy was performed by Parker and Adams [194, 195]. These authors considered the stirring of two-dimensional condensates by a rotating potential, modelled by integrating the time-dependent GPE in the laboratory frame, with a time-dependent rotating potential, using the Crank-Nicolson method (which is a finite-difference scheme; see section 4.3.2). In their simulations, the field appears to become turbulent following the onset of the instability, and they identified the characteristic *Kolmogorov spectrum* of turbulence [272] in the nonequilibrium field. However, contrary to the earlier work of Feder *et al.* [80] and Lundh *et al.* [171], they observed the vortex configuration to ultimately relax to a rigid, crystalline lattice. Parker and Adams stressed the role of the breaking of the initial twofold rotational symmetry of the GP solution as being essential in allowing the field to relax to the equilibrium lattice configuration.

They also considered the effect of imposing a momentum space ‘cutoff’ on their simulations, which they implemented by periodically removing the portion of the GP solution outside some radius $\hbar k_c$ in momentum space⁵. The authors interpreted this process as analogous to an evaporative-cooling type procedure, and noted that despite a reduced final atom number in the presence of this elimination procedure, the final vortex configuration and damping rates were insensitive to the cutoff wave vector k_c . They therefore suggested that the relaxation process is a zero-temperature phenomenon.

8.3 Outstanding questions

The description of the condensate-stirring process has presented a strong challenge for many-body theories: to describe the full dynamical evolution of the system, from cold condensate, through strongly turbulent nonequilibrium dynamics, to a rotating thermal equilibrium. Although the work of Kasamatsu *et al.* [145] established the role of the dynamical instability in generating the thermality required to facilitate the nucleation of vortices, and the work of Lobo *et al.* [170] showed that the requisite dissipative dynamics of vortex-lattice formation can emerge from the deterministic dynamics of a nonlinear-Schrödinger classical-field evolution, a number of important questions about the lattice-formation process remained open.

⁵This operation should be contrasted to the deleterious implementations of the projection operator in formally projected classical-field methods (see, e.g., [189]), in which the deletion is performed *continuously*, so that no loss of population actually occurs (see, e.g., [63]).

Accuracy of classical-field approaches

The calculations of [80, 171, 170, 194, 195] each exhibit the nucleation of vortices from the GP equation, while the simulations of [170] and [194, 195] additionally result in the relaxation to a vortex-lattice configuration. Artifacts, inaccuracies, and numerical noise in the solutions of *nonlinear* differential equations such as the GP equation can lead to the solution taking a completely spurious path, and thus produce unphysical results. It is important to investigate the nucleation and damping phenomena exhibited by these earlier works with the more sophisticated projected-GPE approach, to determine whether the outcomes of these previous numerical studies result from the intrinsic nonlinear dynamics of the classical-field solutions, or if they are due to inaccuracies of the numerical methods⁶.

Emergence of thermal behaviour

In the classical-field viewpoint of Lobo *et al.* [170], it is supposed that the nonlinear dynamics of the classical field lead to an *emergent* thermal component of the field, which provides the dissipation required for vortex nucleation and damping. The dynamical creation, evolution and development of the thermal field is a theoretical question of fundamental interest. In the previous calculations, only Lobo *et al.* [170] attempted any characterisation of the thermal field, and they only considered the final equilibrium thermal field, and only quantified its temperature. It is therefore important to perform a quantitative and detailed examination of the *development* of the thermal component of the field following the onset of the dynamical instability of the quadrupole oscillations of the condensate.

Role of symmetry breaking

The authors of [195] stress the importance of the breaking of the twofold rotational symmetry of the system in the processes leading to lattice formation. In all previous theoretical studies, the rotational symmetry of the condensate was broken either by numerical error [80, 171, 170], by *ad hoc* seeding of field modes with noise [200], or by including perturbations of the trapping potential [145, 194, 195] intended to mimic the noise sources present in realistic experiments. None of these approaches to the symmetry breaking is satisfactory, as they rely either on the explicit failure of the integrator in propagating the classical-field solutions, or relate to physical effects in a somewhat arbitrary and uncontrolled way. The truncated-Wigner theory (section 3.4), on the other hand, informs us that any classical-field trajectory must contain a certain irreducible level of fluctuations,

⁶We follow here a guiding principle of scientific research informally summarised by Feynman [92]: “... if there is something really there, and you can’t see good because the glass is foggy, and you polish the glass and look clearer, then it’s more obvious that it’s there, not less.”

corresponding to the quantum uncertainty present in the corresponding quantum field theory, even at zero temperature. Thus by including this representation of *quantum noise* in the initial conditions of field trajectories, we can investigate the dynamical instability and the resulting thermalisation of the field in a *minimal* stirring scenario, in which the system is considered to be otherwise ‘ideal’, and the only source of symmetry breaking is the inescapable quantum-uncertain nature of the true physical system.

Nature of the equilibrium state in 2D

Finally, a major question raised by the simulations of [80, 171, 194, 195] is the nature of the equilibrium configuration of vortices in 2D systems. Both [80] and [171] reported highly turbulent motion of the vortices at the end of their simulations. The authors of the former paper conjectured that the lattice crystallisation process was dependent on some dissipative process beyond the GP theory, such as coupling to the thermal cloud, while the authors of [171] suggested (in light of the then preliminary results of [170]) that the time scale for relaxation in their 2D system may simply be very long. In contrast, Parker *et al.* [194, 195] have shown in their 2D calculations the relaxation of the turbulent field subsequent to the instability to a rigid vortex lattice, and suggest that the formation is a zero-temperature phenomenon. It is therefore important to determine conclusively the nature of the equilibrium state of the stirred condensate in 2D, and whether the lattice-formation scenario observed in [194, 195] is physical or simply the result of numerical error. This is of particular importance as all experiments on condensate stirring to date have been performed in fully 3D geometries, while two-dimensional models are often assumed in theoretical investigations of vortex dynamics in Bose condensates.

Chapter 9

Condensate stirring

In this chapter, we apply the PGPE formalism to the nonequilibrium dynamics of the condensate stirring scenario. We consider a highly oblate (quasi-2D) condensate, initially at zero temperature in a cylindrically symmetric trap. In accordance with the truncated-Wigner interpretation of classical-field trajectories, we include small Gaussian fluctuations of the classical field, representing the initial vacuum occupation of Bogoliubov modes, and explicitly breaking all symmetries of the Hamiltonian. We then abruptly introduce a weak elliptical trap anisotropy, and simulate the ensuing dynamics of the atomic field in the rotating frame. We observe the rapid thermalisation of material ejected from the condensate bulk during the dynamical instability of the quadrupole mode. This material forms a rotating thermal cloud which drives the Landau instability of surface modes and leads to the nucleation of vortices. We observe the damping of the turbulent vortex motion, and the slow rotational equilibration of the vortices with the rotating trap and thermal cloud. We find, however, that the equilibrium state of the stirred atomic field in this quasi-2D geometry is not a rigid vortex lattice but a disordered *vortex-liquid* state. We characterise the thermodynamic parameters of the thermal field, and monitor the change in these parameters during the equilibration of the system. We find that the final thermodynamic parameters of the field are well predicted by simple arguments based on energy conservation. We discuss the absence of true Bose condensation in the equilibrium vortex-liquid state, and discuss how the temporal correlations of this phase distinguish it from the thermal component of the field.

This chapter is organised as follows. In section 9.1 we define the parameters of the physical system we model, and our classical-field approach to its evolution. In section 9.2 we present the results of a typical, representative simulation. In section 9.3 we present a detailed analysis of this representative classical-field trajectory, quantifying the rotational response and the heating of the field during its evolution, and characterising the field

coherence. In section 9.4, we consider how the behaviour of the field depends on the chemical potential (and thus size and interaction strength) of the initial condensate mode, and the effect of the classical-field cutoff height on the field dynamics. In section 9.5 we summarise our findings and present our conclusions.

9.1 System and simulation procedure

9.1.1 The system

Following chapters 6 and 7, we choose physical parameters corresponding to ^{23}Na atoms confined in a strongly oblate trap, with trapping frequencies $(\omega_r, \omega_z) = 2\pi \times (10, 2000)$ rad/s. We consider an initially zero-temperature, vortex-free condensate characterised by its chemical potential μ_i in this trapping potential. In this chapter we consider condensate chemical potentials $4\hbar\omega_r \leq \mu_i \leq 20\hbar\omega_r$, corresponding to condensate populations $1600 \lesssim N_0 \lesssim 42\,000$. At $t = 0$, the trap is deformed into an ellipse in the xy plane which rotates around the z axis with angular frequency Ω . In the frame co-rotating at frequency Ω , the single-particle Hamiltonian is of the form of equation (3.1), with trap anisotropy

$$\delta V(\mathbf{x}) = \epsilon m \omega_r^2 (y^2 - x^2). \quad (9.1)$$

In this chapter we set the trap eccentricity to $\epsilon = 0.025$, which is sufficient to excite the quadrupole instability at $\Omega \sim \omega_r/\sqrt{2}$ [220, 240]. We will always drive the system within the hydrodynamic resonance, choosing $\Omega \equiv 0.75\omega_r$. The sudden turn-on of the rotating potential at fixed ϵ and Ω means that the classical-field energy (equation (3.8)) of the initial state evaluated in the rotating frame is a conserved quantity of the motion. We again adopt a 2D representation of the system, formalised in our classical-field method by choosing a cutoff $E_R = 30\hbar\omega_r \ll \hbar\omega_z$, such that the low-energy space \mathbf{C} excludes all modes with excitation along the z axis (see section 4.1.1).

9.1.2 Initial state: inclusion of vacuum noise

Symmetry of the zero-temperature GP eigenstate

A key issue in studying the condensate-stirring scenario within the framework of classical-field theory is in identifying the *path* that the classical field takes in transforming from a vortex-free, zero-temperature condensate to a rotating (and, as we will show) finite-temperature field configuration. In general, we expect the field to acquire thermality during its real-time evolution due to the ergodicity of the PGPE evolution, as discussed

in section 3.3. However, we face a significant issue in the present case where we seek to describe the thermalisation of an initially *zero-temperature* condensate. The $T = 0$ ground GP state of the isotropic trap is indeed highly excited with respect to the ground GP state (of the same normalisation) in the rotating frame in which we expect equilibrium to result, and at equilibrium we expect the excess energy of the initial state to be distributed throughout the configuration space of the field in the form of thermal fluctuations. However, as the initial GP state possesses circular rotational symmetry, the exact evolution of the state under the GPE with the elliptic potential we consider would retain unbroken twofold rotational symmetry for all time, and true thermalisation of the field would not occur¹. Moreover, the high degree of symmetry of the GP state impedes the dynamical instability, which results from the exponential *amplification* of deviations about an otherwise stable GP trajectory (section 2.2.3). In the limit of perfect numerical evolution of an initial GP eigenstate under the (projected) GP equation, no such deviations arise, and thus no dynamical instability will occur. In *practical* numerical simulations the build-up of numerical error with time will eventually facilitate the dynamical instability and the ergodic thermalisation of the field. However, this yields somewhat arbitrary results for the time scales for symmetry breaking and the relaxation processes which require it. It is important then to ensure that the symmetry of the classical-field system is broken in a physically defensible way, rather than to rely on uncertain and uncontrolled numerical inaccuracies.

Truncated-Wigner prescription for noise

The resolution of the issue of symmetry breaking is found in the truncated-Wigner representation of the quantum-field theory for the Bose gas discussed in section 3.4. In the truncated-Wigner formalism, the PGPE appears as a classical equation of motion for field trajectories which sample the approximate evolution of the density matrix of the quantum-field system. We recall that in this approach the classical-field trajectories are regarded as stochastic processes, even though in the truncated-Wigner formalism appropriate to the interacting atomic Bose gas, no explicit dynamical noise processes appear. However, a finite amount of noise does appear in the initial conditions, even at zero temperature, representing quantum uncertainty in the physical state. This provides a fundamental and quantitatively correct quantum-mechanical mechanism that breaks the symmetry of the Hamiltonian. While for finite-temperature states this noise can be safely neglected due to

¹In fact, we can identify the highly symmetric GP eigenstate as a member of the ‘set of measure zero’ for which ergodicity does not apply (section 3.3).

the presence of *thermal* fluctuations, at zero temperature many important processes can only be initiated by quantum fluctuations.

Thus, although we consider (as in chapter 7) the relaxation and equilibration dynamics of the classical-field trajectory on long time scales, precluding a formal reconstruction of quantum-field expectation values from the classical-field statistics, we follow the truncated-Wigner prescription in seeding our initial states with the *irreducible* fluctuations of the quantum field theory which are essential for the initiation of the thermalisation dynamics of the field. We therefore add a representation of the vacuum fluctuations to our initial states in the form of classical (complex) Gaussian noise of mean population equal to $\frac{1}{2}$ -quantum per mode. It is necessary, therefore, to choose appropriate basis modes to which the noise is added, so as to faithfully represent the ground state of the many-body system [246]. For the interacting system, the appropriate basis modes in which to construct the vacuum will be distinct from the basis modes in which the system is propagated. In this work our primary aim is to include the possibility of spontaneous processes in the initial condition, so as to initiate the dynamical instability, and facilitate the ergodicity of the classical-field evolution. It therefore suffices to populate the Bogoliubov modes orthogonal to the condensate mode with noise [100, 50, 246]². The initial state is constructed from the GP ground state of the symmetric trap in the lab frame, $\phi_0(\mathbf{x})$, according to

$$\psi(\mathbf{x}) = \phi_0(\mathbf{x}) + \sum_j \alpha_j u_j(\mathbf{x}) + \alpha_j^* v_j^*(\mathbf{x}), \quad (9.2)$$

where (u_j, v_j) are the Bogoliubov modes orthogonal to $\phi_0(\mathbf{x})$, and $\alpha_j = (\xi_j + i\eta_j)/2$ for our system where the thermal population is initially zero. The independent (real) Gaussian distributed variables satisfy $\overline{\xi_j} = \overline{\eta_j} = \overline{\eta_i \xi_j} = 0$ and $\overline{\xi_i \xi_j} = \overline{\eta_i \eta_j} = \delta_{ij}$.

Population of the quasiparticle basis constructed in this manner ensures that all surface modes of the condensate (that are resolvable within the condensate band) are seeded by noise, including those that are dynamically unstable in the presence of the rotating trap anisotropy. The noise introduced here plays a role entirely analogous to that played by vacuum electromagnetic field fluctuations in triggering the spontaneous decay of a two-level atom [230]. All symmetries of the mean-field state are thus broken [236] *before* the exponential growth of noncondensed field density [49] occurs during the dynamical instability.

²Self-consistent prescriptions for the addition of noise to the truncated-Wigner initial state exist [236, 238, 239], but the resulting higher-order corrections to the vacuum state are not important for our application here.

Validity of adding Wigner noise

As we have noted in section 3.3, due to the ergodicity of the PGPE, classical-field simulations tend to equilibria in which the energy of the field is equipartitioned over its normal modes. Consequently, the energy (and population) added to the initial states of classical-field simulations according to the truncated-Wigner prescription eventually thermalises, and contributes spuriously to the equilibrium thermal fraction of the field (section 3.4.3). Sinatra *et al.* [239] thus suggest a validity criterion $T_{\text{class}} - T \lesssim T$ for the (initial) temperature T of the simulation relative to the equilibrium temperature of the classical field T_{class} . We note however that the underlying issue is the relative quantity of energy added in the Wigner noise, as a fraction of the total thermal energy in the field. In our stirring scenario, as viewed in the rotating frame, the initial GP state is a highly *collectively* excited state, the energy of which (above that of the ground vortex lattice state in this frame) is converted into thermal energy as the field relaxes (see section 9.3.1). The energy added to the initial state according to the truncated-Wigner prescription is therefore dwarfed by the energy liberated from the GP state during its relaxation. Similarly, the normalisation and angular momentum of the initial vacuum noise constitute very small fractions of the total quantities contained in the equilibrium thermal field, and so we can be confident that the spurious thermalisation of the vacuum noise does not affect the relaxation dynamics of our trajectories significantly.

9.1.3 Evolution

Using the initial state described above, we evolve field trajectories in the frame of the rotating trap ellipticity, i.e., we apply the projector in the frame $\Omega_p = \Omega = 0.75$. Our procedure corresponds to the sudden introduction of the rotating anisotropy at time $t = 0$. We evolve each trajectory with the ARK45-IP algorithm (section B.2), with accuracy chosen such that the relative change in field normalisation is $\leq 4 \times 10^{-9}$ per time step taken, and the simulations presented here all have total fractional normalisation changes $\lesssim 2 \times 10^{-4}$ over their duration ($\sim 10^4$ trap cycles). Importantly, the change in rotating-frame *energy* is of magnitude commensurate with that of the change in normalisation, so that the truncation error represents a loss of population from the system as a whole, rather than a preferential removal of high-energy components as in [195], and the relaxation of the condensate-band field is thus due to internal damping processes only.

9.2 Simulation results

We present here the results of a simulation with initial chemical potential $\mu_i = 14\hbar\omega_r$ (corresponding to condensate with $N_0 \approx 21\,500$ atoms), whose response to the stirring is representative of systems throughout the range $4\hbar\omega_r \leq \mu_i \leq 20\hbar\omega_r$. Using the criteria discussed in section 3.2.2, we set the condensate-band cutoff $E_R = 3\mu_i = 42\hbar\omega_r$. The upper limit of the range investigated ($\mu_i \sim 20\hbar\omega_r$) is set by computational limitations: at the fixed rotation frequency $\Omega = 0.75\omega_r$ the size of the Gauss-Laguerre (GL) basis scales approximately as $\mathcal{M} \propto E_R^2$, and the corresponding computational load scales as $\mathcal{O}(E_R^3)$ (section 4.2), so that simulations rapidly become numerically expensive. For the cutoff $E_R = 42\hbar\omega_r$ the GL basis consists of $\mathcal{M} = 2028$ modes. At the lower end of the range ($\mu_i \lesssim 10\hbar\omega_r$) the simulation results differ somewhat both in the thermalisation process leading to vortex nucleation, and the behaviour of the vortex array, due to the reduced mean-field effects and small numbers of vortices, respectively. These differences will be discussed in section 9.4.1.

We will focus primarily on the case of $\mu_i = 14\hbar\omega_r$, and note for comparison that a pure condensate (i.e., ground GP eigenstate) with the same number of particles rotating at $\Omega = 0.75\omega_r$ would contain a lattice of ≈ 18 vortices. The response to the stirring is illustrated in the sequence of density distribution plots shown in figure 9.1, and exhibits the following key features.

Dynamical instability and formation of thermal cloud

The initial state (figure 9.1(a)) quickly becomes elongated, with its long axis oscillating irregularly relative to the major axis of the trap. The quadrupole oscillations are dynamically unstable to the stirring perturbation [220, 240, 196], and since all the Bogoliubov modes of the condensate are effectively populated due to the representation of quantum noise in our simulation, some of the fluctuations grow exponentially. The effect is clearly seen in figure 9.1(b), where matter streams off the tips of the elliptically deformed condensate in the direction of the trap rotation. Material is then ejected more or less continuously, forming a ring about the condensate, with the latter oscillating in its motion and shape. The ring soon becomes turbulent and diffuse (figure 9.1(c)), losing coherence and forming a thermal cloud whose characterisation (section 9.3.1) is a central theme of this chapter. In section 9.3.2 we show from an analysis of the angular momentum and density distribution that by $t = 50$ cyc the outer part of the cloud (beyond $r \approx 7r_0$) rotates as a normal fluid in rotational equilibrium with the drive.




Figure 9.1: Classical-field density in the rotating frame at representative times during the evolution. Shown are (a) initial condition: lab-frame ground state plus vacuum noise, (b) ejection of material during the dynamical instability, (c) nucleation of vortices at the interface between the condensate and the noncondensate material, (d) penetration of vortices into the condensate bulk, (e) nonequilibrium state with rapid vortex motion, and (f) equilibrium state. Vortices are indicated by white circles, and are shown only where the surrounding density of the fluid exceeds some threshold value, and + marks the coordinate origin (trap axis). Parameters of the simulation are given in the text.

Vortex nucleation

After the thermal cloud has formed, surface oscillations of the condensate grow in magnitude (figure 9.1(d)), in accordance with the prediction of Williams *et al.* [266], that such oscillations are unstable in the presence of a thermal cloud. These fluctuations form a transition region between the central condensate bulk, characterised by comparatively smooth density and phase distributions (see also figure 9.8(a)), and an outer region of thermally occupied highly energetic excitations. Long-lived vortices are nucleated in this transition region, and then begin to penetrate into the edges of the condensate region. Initially the vortices penetrate only small distances $r \approx r_0 = \sqrt{\hbar/m\omega_r}$ with rapid and irregular motion, and visibly lag the rotation of the drive. This process is one in which the *emergent* thermal component of the field acts to provide the dissipation which was included phenomenologically in the simulations of Tsubota *et al.* [254] and Penckwitt *et al.* [200].

Vortex-array formation

The nucleated vortices remain near the edge of the condensate for a considerable time before they begin to penetrate substantially into the condensate. Eventually vortices pass through the central region (by $t \approx 75$ cyc in this simulation), with trajectories that are largely independent, except when two vortices closely approach each other³. The density of vortices within the central region gradually increases (figure 9.1(e)), and the vortices begin to form a rather disorderly assembly, which still lags behind the drive in its overall rotation about the trap axis. The erratic motion of individual vortices and their differential rotation with the respect to the drive then gradually slow, and they become more localised, until by $t \approx 150$ cyc they have formed into an array rotating at the same speed as the drive. This formation process can be interpreted as the damping of vortex motion by the mutual friction between them and the rotating thermal cloud [243, 75]. Alternatively, we can interpret the motion of the vortices as the result of low-energy excitations of the quantum fluid, which are gradually damped by their interaction with high energy excitations [82, 112, 206]. The coupling of the high-energy modes to the low-energy excitations of the fluid drives their distribution towards equilibrium [225, 237]. The high energy modes comprising the thermal cloud equilibrate to a distribution with identifiable *effective* thermodynamic parameters (see section 9.3.1) on a relatively short time scale (~ 100 cyc). While the effective temperature remains approximately fixed, the effective chemical potential subsequently increases over a longer time scale (~ 3000 cyc), as the

³In this and other simulations we have observed a tendency of vortices to cluster together in (short-lived) co-rotating pairs which orbit their geometric centre, a behaviour characteristic of two-dimensional superfluid turbulence (see reference [259] and references therein).

distribution is shaped by its interaction with the low energy excitations. By $t \approx 1000$ cyc the vortex motion (as quantified by the mean vortex speed; section 9.3.4) reaches its equilibrium level, although the *density* of vortices in the central bulk of the field continues to grow until $t \approx 5000$ cyc (section 9.3.2). We find that the equilibrium populations of thermal excitations in this final state are such that the vortex distribution does not ‘crystallise’ into a rigid lattice, and is more appropriately characterised as an equilibrium *vortex liquid*, as we discuss in section 9.3.3.

9.3 Analysis

In this section we present the set of measurements we use to analyse the properties of the thermal cloud, and to characterise the coherence of the central region in which the vortices reside. We begin by making a simple estimation of the energy and temperature the thermal cloud would acquire during an idealised lattice formation scenario. Next we obtain the chemical potential of the thermal cloud by using a self consistent fitting procedure for the atomic distribution function, which also provides a measure of the temperature. We characterise the rotational properties of different spatial regions of the system, which provides evidence that during the vortex nucleation stage, the central part behaves as a superfluid, while the outer cloud behaves as a classical gas, rotating like a rigid body.

The issue of condensate identification is a very important one, and is most commonly done in terms of spatial correlation functions, using ensemble averages of quantum mechanical operators [164, 60]. However, we will show that the varying vortex configurations within the ensemble have a destructive effect on the spatial order of the system we study here. This leads us, in section 9.3.3, to investigate temporal correlation functions, which we show can be used to characterise the local coherence of the field, and consequently distinguish the quasi-coherent behaviour of the central region from that of the thermal cloud.

Following the first appearance of an identifiable vortex array, there is a long period over which it slowly relaxes to an equilibrium state. In section 9.3.4 we measure the motional damping of the vortices and interpret it in terms of their interaction with the thermal component of the field [81]. Note that all results and figures in this section (figures 9.2-9.10) correspond to the simulation shown in figure 9.1.

9.3.1 Thermodynamic parameters

Analytic predictions

Since our system conserves both normalisation and energy in the rotating frame, some simple analytic predictions can be made about the development of the thermal component of the field, using the Thomas-Fermi (TF) approximation. In a frame rotating at angular velocity Ω , the ground state is a vortex lattice. In the Thomas-Fermi approximation, the vortices are replaced by a centrifugal dilation of the trapping potential arising from the coarse-grained (rigid-body) flow pattern, yielding the TF wave function

$$\psi_{\text{TF}}^{\Omega} = \sqrt{\frac{\mu_{\Omega} - m(\omega_r^2 - \Omega^2)r^2/2}{U_{2\text{D}}}} \Theta\left(\mu_{\Omega} - m(\omega_r^2 - \Omega^2)\frac{r^2}{2}\right), \quad (9.3)$$

where μ_{Ω} is the chemical potential of the rotating frame solution. From equation (9.3) we obtain the number of particles N_{TF}^{Ω} and energy E_{TF}^{Ω} ⁴ for the lattice state

$$N_{\text{TF}}^{\Omega} = \int d\mathbf{x} |\psi_{\text{TF}}^{\Omega}|^2 = \frac{\pi\mu_{\Omega}^2}{U_{2\text{D}}m(\omega_r^2 - \Omega^2)}, \quad (9.4)$$

$$E_{\text{TF}}^{\Omega} = \int d\mathbf{x} \left[\frac{m}{2}(\omega_r^2 - \Omega^2) + \frac{U_{2\text{D}}}{2} |\psi_{\text{TF}}^{\Omega}|^2 \right] |\psi_{\text{TF}}^{\Omega}|^2 = \frac{2}{3}\mu_{\Omega}N_{\text{TF}}^{\Omega}. \quad (9.5)$$

The corresponding quantities for the vortex-free inertial-frame ground state are given by setting $\Omega = 0$. We note that the chemical potential $\mu_0 \equiv \mu_{\Omega=0}$ of the nonrotating state has the same value in both the inertial and rotating frames. If we now compare a ‘lattice’ state (i.e. $\psi_{\text{TF}}^{\Omega}$) and a vortex-free state (ψ_{TF}^0), both fully condensed and with equal occupation, we see from equation (9.4) that the chemical potentials are related by

$$\mu_{\Omega} = \mu_0 \sqrt{1 - \Omega^2/\omega_r^2}, \quad (9.6)$$

and hence from equation (9.5) that their rotating-frame energies are related by

$$E_{\text{TF}}^{\Omega} = E_{\text{TF}}^0 \sqrt{1 - \Omega^2/\omega_r^2}. \quad (9.7)$$

The excess rotating-frame energy of the vortex-free ground state over the lattice state of the same number of atoms,

$$\Delta E \equiv E_{\text{TF}}^0 - E_{\text{TF}}^{\Omega} \quad (9.8)$$

can be significant, and for example is approximately $E_{\text{TF}}^0/3$ for an angular velocity of

⁴The appropriate Thomas-Fermi energy functional is similarly given by equation (3.8) upon neglecting the kinetic energy, and including the appropriate centrifugal dilation of the trapping potential.

$\Omega = 0.75\omega_r$. Thus in our stirring scenario, the rotating equilibrium condensate state reached at the end of the process must have less energy and less atoms than the initial vortex-free state, and the excess energy and atoms constitute a thermal cloud⁵. The exact result for the excess energy of the classical field, obtained from the simulations, will depend upon the depletion of the condensate mode required to form the thermal cloud and the mutual interaction of the cloud with the condensate. For small thermal fractions this energy could be found by means of a calculation similar to that described in reference [236]. In the simplest approximation we assume the limit of a small thermal fraction, in which the energies of the condensate and thermal field are additive, with equation (9.8) approximating the excess energy the thermal field contains. We further assume that in equilibrium this energy will be classically equipartitioned (section 3.3.2) over weakly interacting harmonic oscillators, in the spirit of the Bogoliubov approximation [170]. The equilibrium temperature can thus be predicted by

$$k_{\text{B}}T = \frac{\Delta E}{\mathcal{M} - 1}, \quad (9.9)$$

with \mathcal{M} the condensate-band multiplicity. For a given initial chemical potential we see therefore that the equilibrium temperature will be strongly dependent on the basis size. However, we shall see in section 9.4, equation (9.6) provides a very good estimate of the final chemical potential reached by the field, which is essentially independent of cutoff.

Self-consistent fitting

We now estimate the thermodynamic parameters of the nonequilibrium classical field using the semiclassical fitting-function approach of section 6.5.4. We remind the reader that in this approach we calculate the profile of the thermal cloud in the Hartree-Fock potential due to the *entire* classical field, and thus do not need to explicitly distinguish the condensed component of the cloud from the thermal component in forming the mean-field potential. This is particularly advantageous in the present scenario as no true condensate exists in the turbulent field, as we discuss in section 9.3.3.

Our task is somewhat simplified in the present case, as the thermal cloud rapidly comes to equilibrium with the rotating drive (as we show in section 9.3.2), and so we assume that the cloud rotation rate⁶ $\Omega_{\text{th}} = \Omega$. The fitting function for the radial density

⁵This energy change ΔE cannot be attributed to the initial trap deformation associated with the stirrer, as the energy of an axially symmetric state is unchanged by the abrupt introduction of a perturbation of quadrupolar symmetry (i.e., within the sudden approximation [39] implicit in our procedure).

⁶As noted in section 6.5.4, the results of the fitting procedure are rather insensitive to the precise value of angular velocity used.




Figure 9.2: Fitting procedure for the thermodynamic parameters. Shown are: the (time-averaged) radial density distribution times radius (solid line), the effective potential experienced by noncondensate atoms (dashed line), and the fitted distribution of noncondensed atoms times radius (dash-dotted line). The dotted line indicates the classical turning point corresponding to energy cutoff E_R . The data shown corresponds to the period $t = 3000\text{--}3010$ cyc (see figure 9.1(f)).

distribution thus reduces to (appendix D)

$$n_{\text{th}}^{\text{fit}}(r; \mu, T) = \frac{1}{\lambda_{\text{dB}}^2} \ln \left[\frac{E_R - \mu}{\langle V_{\text{eff}}^{\text{HF}}(r) \rangle - \mu} \right], \quad (9.10)$$

where the effective potential

$$\langle V_{\text{HF}}^{\text{eff}}(r) \rangle = m (\omega_r^2 - \Omega^2) \frac{r^2}{2} + 2U_{2\text{D}} \langle \tilde{n}(r) \rangle, \quad (9.11)$$

with $\langle \tilde{n}(r) \rangle$ the azimuthally averaged density of the field (section 6.5.4), and we present an example of such a fit in figure 9.2.

After the initial strongly nonequilibrium dynamics following the dynamical instability (i.e., by $t \approx 400$ cyc), the averaged field densities $\langle \tilde{n}(r) \rangle$ are well fitted by the function given in equation (9.10), and we conclude that the higher energy components of the field have reached a quasistatic equilibrium. We then follow the evolution of the temperature and chemical potential with time, which we present in figure 9.3. The temperature (figure 9.3(a)) appears to be essentially constant within the accuracy of the measurement procedure, while the chemical potential (figure 9.3(b)) shows a definite upward trend over




Figure 9.3: Evolution of the effective (a) temperature and (b) chemical potential of the noncondensed atoms, as measured by the semiclassical fitting procedure.

the course of the field's evolution. In section 9.3.3 we compare the chemical potential of the cloud determined from our fitting procedure with an effective chemical potential of the dense central vortex-containing region, which we estimate by considering the temporal frequency components present in the field.

9.3.2 Rotational parameters

As noted in section 6.4, a rotating superfluid exhibits a moment of inertia which is suppressed below its classical value. As a superfluid is rotated at an increasing angular velocity, it admits an increasing number of vortices, and its moment of inertia tends towards the classical value [128]. In particular, as the rotation frequency Ω of a zero-temperature condensate (GP state) increases, the (areal) density of vortices increases [99, 79], and is given to leading order by the Feynman relation (section 2.3.3)

$$\rho_v^c = \frac{m\Omega}{\pi\hbar}. \quad (9.12)$$

with corrections arising due to the lattice inhomogeneity and the discrete nature of the vortex array [233]. As the rotation frequency and thus the vortex density increase, both the classical and quantum moments of inertia increase also, and the difference between the two moments vanishes in the limit of rapid rotations $\Omega \rightarrow \omega_r$ [79].

Vortex density

To monitor the increase in vortex density during the system evolution, we count the vortices of positive rotation sense occurring within a circular counting region of radius equal to the TF radius of the initial state, $R_{\text{TF}} = \sqrt{2\mu_i/m\omega_r^2}$. Due to the centrifugal dilation of the rotating condensate this region lies within the central bulk of the condensate, and in this manner we attempt to avoid including in the count the short-lived vortices constantly created and annihilated at the condensate boundary. At equally spaced times over a period of 4 trap cycles about time t_i we obtain the average number of vortices in the counting radius, $\overline{n_v(t_i)}$, and from this we determine the average vortex density in the counting region

$$\rho_v(t_i) = \frac{\overline{n_v(t_i)}}{\pi R_{\text{TF}}^2}. \quad (9.13)$$

In figure 9.4(a) we plot equation (9.13) as a fraction of the prediction of equation (9.12). During the dynamical instability and initial nucleation of vortices the vortex densities measured in this manner are spuriously high, due to the counting of phase defects in the turbulent nonequilibrium fluid rather than long-lived vortices in the condensate bulk. After this initial period, i.e. from $t \approx 150$ cyc onwards, we see a gradual increase in ρ_v as the condensate slowly relaxes and admits more vortices into its interior, and approaches rotational equilibrium with the trap. By $t \approx 5000$ cyc (inset to figure 9.4(a)), the vortex density appears to have essentially saturated at a level $\rho_v \approx 0.85\rho_v^c$. Such a value seems reasonable for the rotation rate and condensate size considered here, which are such that inhomogeneity and discreteness effects will be significant [233]⁷.

Local rotational properties

In section 6.4 we presented an analysis of the rotational properties of the condensate and thermal field in a rotating system, based on a separation of the one-body density matrix into condensed and noncondensed parts. As we discuss in section 9.3.3, the irregular motion of vortices throughout the field we consider in this chapter suppresses the spatial order of the system and thus destroys the condensation in the field. In order to characterise the rotation of the outer cloud formed from ejected material as well as that of the central bulk of the field we thus measure *localised* expectation values, i.e. expectation values of the classical field evaluated over a restricted spatial region. In this way we can compare the mass distribution and angular momentum of particular regions, in order to characterise

⁷Moreover, as noted by the authors of [233], their idealised $T = 0$ model neglects all thermal excitations. It is therefore likely that thermal effects will further affect the vortex density in a finite-temperature vortex lattice and *a fortiori* the vortex liquid we consider here.

the localised properties of the field. We define the expectation value of an operator \mathcal{O} on spatial domain \mathcal{D} by

$$\langle \mathcal{O} \rangle_{\mathcal{D}} = \int_{\mathcal{D}} d\mathbf{x} \psi^* \mathcal{O} \psi, \quad (9.14)$$

where the \mathcal{D} will be either region \mathcal{A} or \mathcal{B} illustrated in figure 9.5. Focussing our attention on the outer annulus \mathcal{B} , we calculate the classical moment of inertia of material in the annulus

$$\langle \Theta_c \rangle_{\mathcal{B}} = m \langle r^2 \rangle_{\mathcal{B}}, \quad (9.15)$$

and the angular momentum of this material $\langle L_z \rangle_{\mathcal{B}}$. From measurements over the period $t = 50 - 51$ cyc we find that

$$\frac{\langle L_z \rangle_{\mathcal{B}}}{\Omega} = \langle \Theta_c \rangle_{\mathcal{B}}, \quad (9.16)$$

to within 4%, indicating that the cloud in this region is rotating as a normal fluid in rotational equilibrium with the drive. Averaging samples of these expectation values over a period of a trap cycle (which we denote by an overline $\overline{\cdots}$), we find for the time averages $\overline{\langle L_z \rangle_{\mathcal{B}}}/\Omega = 1.025 \overline{\langle \Theta_c \rangle_{\mathcal{B}}}$. Turning our attention to the central disc \mathcal{A} , we find that over the same period, $\overline{\langle L_z \rangle_{\mathcal{A}}}/\Omega = 0.080 \overline{\langle \Theta_c \rangle_{\mathcal{A}}}$. The central condensate bulk at this time thus remains approximately stable against vortex nucleation, but possesses some small angular momentum due to its shape oscillations, in contrast to the ejected material which has lost its superfluid character and has come to rotational equilibrium with the drive.

Global rotational properties

We consider now the rotational properties of the entire field, by evaluating expectation values as above over the full xy plane (i.e., in equation (9.14) we set $\mathcal{D} \rightarrow \mathbb{R}^2$). An obvious measure of the field's rotation is given by the angular momentum L_z , which we here scale by Ω to form the quantity Θ (equation (6.16)) which, at equilibrium, corresponds to the (quantum) moment of inertia. We consider the evolution of this quantity with time, and compare it to the classical moment of inertia $\langle \Theta_c \rangle$ (equation (6.15)) which characterises the dilation of the field's mass distribution during the stirring process. In figure 9.4(b) we plot the evolution of these quantities as time evolves and in figure 9.4(c) we plot the evolution of their ratio. Comparing the evolution of these quantities and that of the field's density distribution (cf. figure 9.1), we identify three roughly distinct phases of the system evolution: (i) Excitation of unstable surface mode oscillations accompanied by permanent increase in angular momentum (figure 9.1(b)), (ii) further increase in angular momentum and $\Theta/\langle \Theta_c \rangle$ as vortices are nucleated into the condensate bulk (figures 9.1(c-e)), and (iii) gradual approach of the field to rotational equilibrium with the drive (figure 9.1(f)). Note that although the finite-temperature equilibria here contain a significant normal-fluid




Figure 9.4: Rotational response of the field. Shown are: (a) Vortex density, with dash-dotted lines indicating $\rho_v / \rho_v^c = 0.7$ and $\rho_v / \rho_v^c = 1$ for reference, (b) classical (dashed line) and quantum moments of inertia, and (c) ratio of quantum and classical moments of inertia. Insets show the behaviour of the same quantities over a longer time scale. Note that the initial $\langle L_z \rangle$ is finite due to the initial vacuum occupation. For reference dotted lines in (a-c) separate phases of the evolution discussed in the text.




Figure 9.5: Regions used for the evaluation of angular momenta and moments of inertia. Disc \mathcal{A} lies entirely within the condensate, where the flow is (initially) irrotational. Annulus \mathcal{B} (between the solid black line and the dashed line indicating the classical turning point of the condensate band) contains only noncondensed, normal fluid. The field density shown is that at $t = 50$ cyc.

component, the quantum moment of inertia is still suppressed below the classical value due to the presence of the superfluid component.

9.3.3 Temporal analysis

Issues in characterising field coherence

A major theme of this thesis is the characterisation of correlations in classical-field trajectories, and the identification of condensation in the classical field. Fundamentally, the correlation functions used to characterise coherence in the formal many-body theory are *ensemble* averages; indeed, these are the only quantities that quantum mechanics (in its most widely accepted interpretations) presumes to describe [39]. Thus, by analogy, we expect the statistics of the classical field to be characterised by ensemble averages over classical-field trajectories. As we have discussed in section 3.3, in thermal equilibrium, we can substitute time averages along single classical-field trajectories for averages over multiple trajectories, exploiting the ergodic hypothesis. Implicit in this approach is the fact that the correlations of an equilibrium field are invariant with respect to displacements in time.

In nonequilibrium scenarios, the correlations of the field are in general time-dependent, and so the situation is more subtle. In chapter 7, we followed [29] in approximating the time-dependent correlations of a *nonequilibrium* field by taking *short-time* averages over the classical field. This approach relies on the system being in some sense *quasi-equilibrium*, so that sufficient statistics to characterise the correlations of interest (i.e., those describing condensation), can be accumulated on a time scale short compared to the ‘macroscopic’ relaxation of the field. However, in taking this approach we must bear in mind that there is no guarantee that the values of correlation functions we measure in this way are representative of the values of the *same* correlation functions evaluated over the appropriate ensemble. In particular, as noted by Leggett [164], one can imagine scenarios in which the trajectories of the system are wildly different, despite each individually exhibiting the coherence properties we might associate with Bose condensation. In such a case we may obtain a large condensate fraction from short-time averages of the individual trajectories, while averages over the formal ensemble reveal no condensate under the Penrose-Onsager definition (section 2.2.1). The scenario, considered in chapter 6, of a vortex precessing at equilibrium is similar: while the characterisation of the coherent condensate from short-time averages in the appropriate rotating frame is unambiguous, this condensate mode does not appear in the corresponding one-body density matrix evaluated over the microcanonical ensemble. The issue is again one of the definition of condensation: the Penrose-Onsager (PO) definition is in some sense the simplest possible measure of the correlations we expect to accompany the presence of condensation in the field, and general scenarios may require the evaluation of more complicated correlation functions in order to appropriately define the ‘condensate’ (see, e.g., references [203, 98, 76, 271]). In general, we expect the appearance of one-body coherence in time averages of single trajectories to correspond to more complicated correlations evaluated over the ensemble. In this sense, measuring the single-trajectory correlations over short time periods constitutes a kind of short-cut to reveal the coherence effects we are ultimately interested in, without having to develop the appropriate ensemble correlation functions and evaluate them formally.

In this chapter, we consider a scenario in which the classical field contains many, irregularly moving vortices, even at equilibrium. It is intuitively clear that averaging over distinct, irregular vortex configurations will yield a suppressed measure of one-body coherence, as for example quantified by the PO approach. Similarly, time averages over a single trajectory of the classical field will also show a suppressed condensate fraction, due to the thermal motion of the vortices over time. Bradley *et al.* [37], considering a similar scenario of thermally excited vortex matter, have shown that a measure of condensation can be obtained by considering short-time averages of the classical field, similarly to those considered in section 7.2.1. It is clear that this approach gives a useful quantification of

the coherence in the system, as evidenced by the behaviour of this condensate fraction as a function of the field temperature, and the results for the growth of coherence in the system with time following a quench of the thermodynamic parameters. However, it is less clear how the coherence measured by this approach corresponds to formal correlations of the many-body system: in general, the condensate fraction one measures in such an approach will depend on the time over which the averaging is performed. In chapter 6, the field coherence was destroyed by the diffusion of the vortex phase resulting from the symmetry breaking; thus the appropriate time scale for averaging was reasonably well defined. In more general situations, however, the ‘natural’ time scale for averaging may not be clear, and the results will depend strongly on the averaging time, which essentially determines the range of temporal frequency (or *energy*) components which contribute to the correlation function⁸.

We will show in this section that the irregular motion of vortices in the final state of the field completely destroys condensation as defined by the PO procedure. Accordingly, there is also no coherently phase-rotating ‘mean field’ as we defined in chapter 5. Nevertheless, we will show that the temporal correlations of the field do give a useful characterisation of the field, and allow us to quantify the time scales over which the field is *locally* coherent. This allows us to distinguish qualitatively distinct phases of the atomic field in both equilibrium and nonequilibrium conditions.

Covariance matrix

We begin by considering the covariance matrix equation (3.26), which we form from 2501 samples of the classical field, taken over the period $t \in [9900, 10\,000]$ cyc. From the analysis of section 9.3.2, we conclude that the trap anisotropy has brought the field to rotational equilibrium in the co-rotating frame, and so we simply construct $G(\mathbf{x}, \mathbf{x}')$ in this frame. Diagonalising this matrix, we find there is no single mode with a large occupation, but rather, the occupations n_i of the ~ 10 most highly occupied eigenmodes $\chi_i(\mathbf{x})$ of $G(\mathbf{x}, \mathbf{x}')$ are comparable, and the occupations of these eigenmodes decay away smoothly with increasing mode number, as shown in figure 9.6(a). As in section 6.3.1, we consider the coherence functions $g_i^{(2)}$ of these modes (inset to figure 9.6(a)). In contrast to the case considered in that section, however, we find that these coherence functions indicate incoherent statistics for all modes (even the ‘most coherent’ highly occupied modes have $g_i^{(2)} > 1.5$, i.e., closer to thermal statistics than coherent statistics). We thus conclude

⁸Note that in the extreme limit that only one sample of the classical field is taken, motion of all frequencies is excluded, and (e.g.) the PO procedure trivially yields a single coherent condensate mode (namely the instantaneous configuration of the entire classical field).

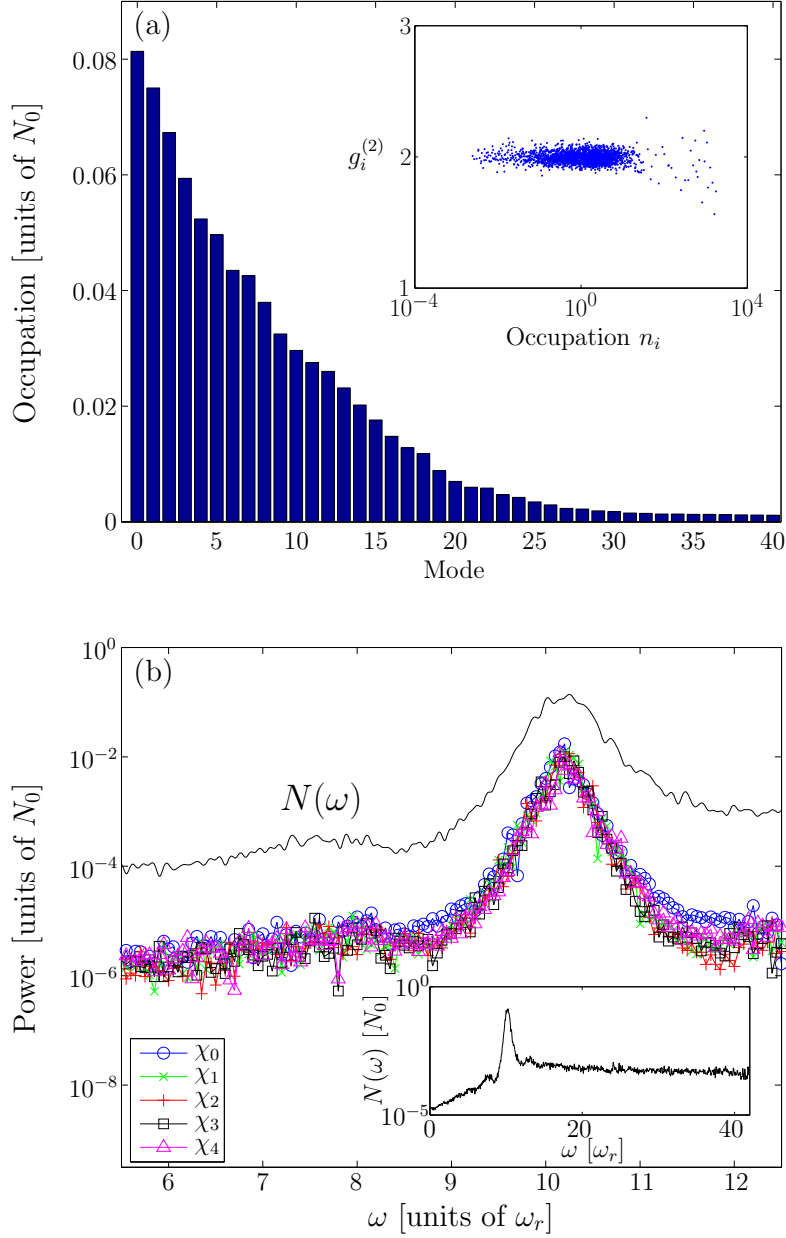


Figure 9.6: (a) Mean occupations of covariance-matrix eigenmodes. Inset: Second-order coherence functions of the modes versus mode occupation. (b) Power spectra of the first five covariance-matrix eigenmodes (lines with symbols). Also indicated is the spatially integrated power spectrum $N(\omega)$ of the whole classical field. The inset shows the behaviour of $N(\omega)$ over a greater range of frequencies.

that the state of the classical field is not that of a fragmented condensate, but rather that there is no coherent condensate mode at all.

Following section 6.3.1, we proceed to calculate the power spectra of the most highly occupied eigenmodes $\chi_i(\mathbf{x}) : 0 \leq i \leq 4$, which we present in figure 9.6(b). We find that all these modes exhibit a power-spectrum peak at $\omega \approx 10\omega_r$. However, in contrast to the narrow peaks of the power spectra in figure 6.5(b), the peaks in this case are very broad (note the logarithmic scale of figure 9.6(b)). The breadth of these peaks is perhaps not surprising, given that the modes have essentially incoherent fluctuation statistics. Nevertheless, the fact that the power in these modes is mostly concentrated in a small frequency range suggests that coherence effects are still important in the field dynamics. For comparison, we calculate also the spatially integrated power spectrum $N(\omega)$ of the entire classical field (as defined in section 5.2.1), calculated here from 1401 samples of the classical field taken over the period 9900 – 9910 cyc. We find that this function also exhibits a broad peak at around the same frequency range as the peaks in the spectra of the individual modes $\chi_i(\mathbf{x})$. We note that the width of this function is significantly greater than that characteristic of the ‘windowing’ effect of the finite sample time, $\Delta\omega = 1/T_{\text{samp}} \approx 0.016\omega_r$, in contrast to the power-spectrum peak of the true condensate mode presented in figure 5.1(b), for which the width was limited by $\Delta\omega$. However, considering this function over a greater range of frequencies (inset to figure 9.6(b)), we see that the behaviour of $N(\omega)$ is broadly similar to that of the true condensate studied in section 5.2.1, with a wide wing of power contained in frequencies greater than that of the peak, with a smaller wing occurring also at frequencies ‘negative’ with respect to the peak. Although no condensate is present in the field, this structure is not unexpected as, for example, Mora and Castin have shown [182] that a Bogoliubov-like picture can be constructed for the description of *quasicondensates* [213], in which density fluctuations are suppressed but true condensation is destroyed by phase fluctuations [205].

Local spectrum of the field

Here we follow section 5.2.1 in analysing the frequencies present in the field at particular points in position space. In order to compare the behaviour of the system at different stages of its nonequilibrium evolution, we define the power spectrum of the classical field ψ at position \mathbf{x} about time t_0

$$H(\mathbf{x}, \omega; t_0) = \left| \mathfrak{F}_{t_0}^T \{ \psi(\mathbf{x}, t) \} \right|^2, \quad (9.17)$$

where $\mathfrak{F}_{t_0}^T\{f(t)\}$ denotes the Fourier coefficient

$$\mathfrak{F}_{t_0}^T\{f(t)\} \equiv \frac{1}{T} \int_{t_0-T/2}^{t_0+T/2} f(t) e^{i\omega t} dt. \quad (9.18)$$

We choose a sampling period T of 10 trap cycles, which is long compared to the time scale $\tau_1 = 1/\omega_p \sim 0.1$ cyc corresponding to the location of the broad peaks in figure 9.6(b) which characterise the phase evolution of the dense central region, while being short compared to that of the relaxation of the field ($\tau_R \sim 1000$ cyc), and choose the sampling interval so as to resolve all frequencies present due to the combined single-particle and mean-field evolutions ($\hbar\omega_{\max} = E_R + 2\mu_i$). As in section 5.2.1, we average this result over azimuthal angle θ , to form $\tilde{H}(r, \omega; t_0)$. In order to gauge the relative strength of the various frequency components at different radii within the classical region, we normalise the resulting data so that the total power at each radius r_j is the same.

In figure 9.7(a), we plot the result of such a calculation over the time interval $t = 9900$ – 9910 trap cycles. For comparison we have overlaid the centrifugally dilated potential $V_{\text{rot}}(r) = m(\omega_r^2 - \Omega^2)r^2/2$ that would be seen by a classical particle in the rotating frame, the cutoff frequency E_R/\hbar , and the semiclassical turning point r_{tp} . The most immediately noticeable feature of this plot is a prominent peak in the power spectrum at $\omega \approx 10\omega_r$, which occurs for radii $r \lesssim 8$, and obviously corresponds to the peaks identified in figure 9.6(b). Despite there being no true condensate in the field, the appearance of this peak clearly indicates that the state of the field in the trap centre is not simply thermal. This result is consistent with analytical results obtained by Graham [118], which suggest that (quasi-)long-range spatial order of the Bose field is accompanied by (quasi-)long-range *temporal* correlations which decay in a functionally equivalent way. We thus associate this frequency with an *effective* (quasi-)condensate eigenvalue $\lambda_e \approx 10\omega_r$. Similarly to figure 5.1, the spectrum shows distortion in the central region where this quasi-coherent object exists, with power occurring at negative frequencies down to $\omega \approx \lambda_e - E_R/\hbar$. In figure 9.7(b) where the power spectrum is plotted at 3 particular radii, the peak in the spectrum obtained close to the trap centre (cyan (light grey) line) can be seen clearly. In addition we see that frequencies above and below this are present, and we identify this with the restructuring of the excitation spectrum of the trap due to the presence of the interacting quasicondensate-like structure (cf. [182]). We note that a secondary peak is visible at $\omega \approx 52\omega_r = \lambda_e + E_R/\hbar$, and an examination of the data reveals that a smaller peak is also present at $\omega \approx -32\omega_r = \lambda_e - E_R/\hbar$. This is an inevitable artifact of the projected method which occurs as follows: the \mathcal{M} Bogoliubov modes which diagonalise the Bogoliubov Hamiltonian in the condensate band in the presence of a stationary condensate

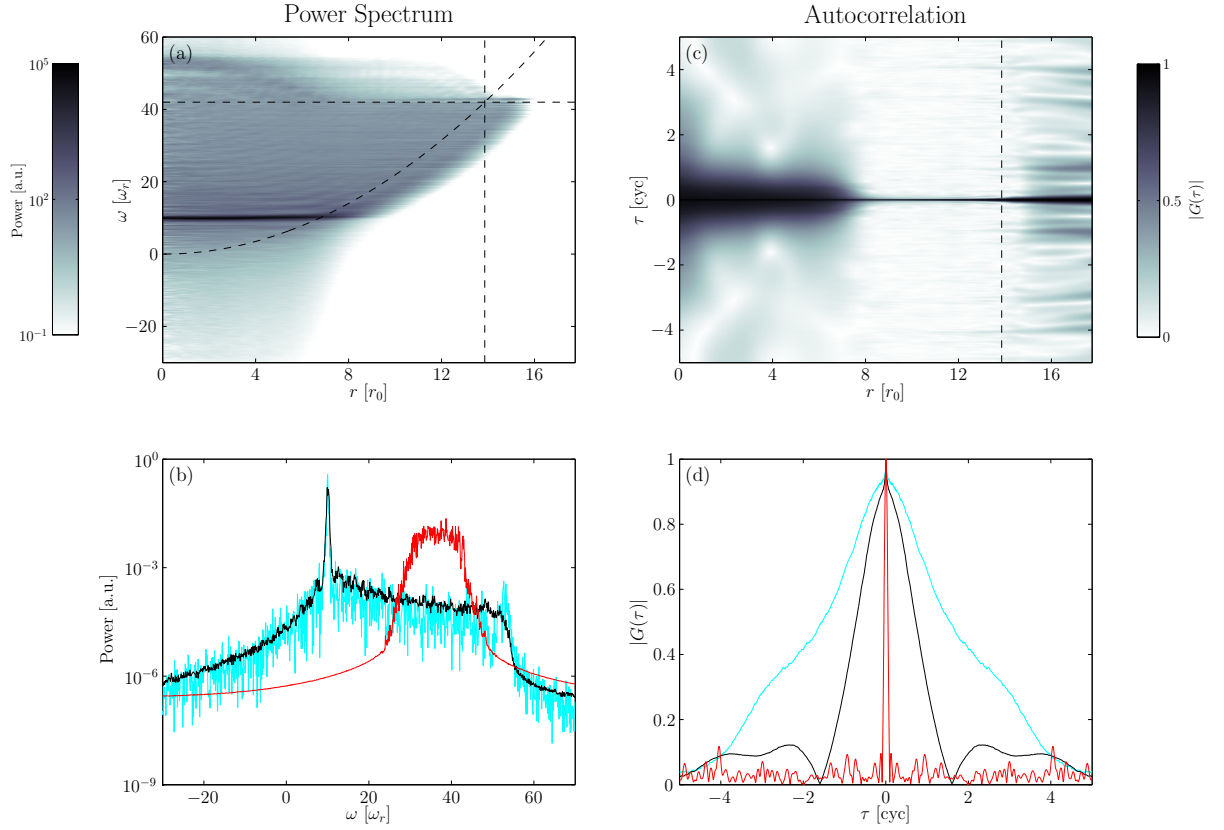


Figure 9.7: (a) Azimuthally averaged power spectral density (PSD). Dashed lines indicate the cutoff energy, corresponding classical turning point and (centrifugally dilated) trapping potential. (b) PSD traces at particular radii. The black and red (dark grey) lines correspond to radii $r = 3.1894r_0$ and $r = 11.9575r_0$, respectively. The cyan (light grey) line corresponds to the smallest radius on the quadrature grid (section C.3), $r_{\min} = 0.0661r_0$. (c) Autocorrelation function (absolute value) obtained from the power spectrum. The dashed line indicates the classical turning point of the condensate band. (d) Autocorrelation magnitude at radii corresponding to figure 9.7(b). Data corresponds to the period $t = 9900 - 9910$ cyc.

mode (GP eigensolution) can be considered as variational approximations to the lowest-lying ‘true’ Bogoliubov modes (obtained as the cutoff $E_R \rightarrow \infty$). The highest energy modes in the restricted quasiparticle basis are poor representations of the true modes, and, to maintain orthogonality of the set, are spuriously localised to the trap centre, and therefore have spuriously high energies $\epsilon_k + \mu_c \approx E_R + \mu_c$ due to the large mean-field interaction there. Therefore, assuming this behaviour of excitations to hold in a finite-temperature equilibrium of the classical field, a peak is expected to occur at these frequencies at the trap centre, where the density of these modes is comparatively large. However, the population contained in the peak observed here is $\sim 0.5\%$ of the total thermal population and as such we do not expect it to qualitatively affect the relaxation process⁹. The black line in figure 9.7(b) indicates the power spectrum at a larger radius $r \approx 3.2r_0$, where the quasicondensate peak is slightly less prominent. At larger radii still ($r \gtrsim 9r_0$) the energy distribution returns to approximately that of the noninteracting gas. As noted in section 5.2.1, frequencies persist into the classically excluded region as a result of the familiar evanescent decay of energy eigenmodes into the potential wall. The effect of the mean-field interaction on the energy spectrum at these large radii, as can be seen in the red (dark grey) line in figure 9.7(b), is much less dramatic (e.g., only frequencies $\omega > \lambda_e$ are present), and the population of frequencies is only significant up to a small increment above the cutoff energy, which we interpret as the mean-field shift of the highest energy single-particle modes. We measure this shift to be $\approx 2.5\%$, indicating that the cutoff we have chosen is sufficiently high to contain all modes significantly modified by the presence of the dense quasi-coherent structure in the field (see section 3.2.2).

Temporal coherence

The temporal analysis of the classical field presented above clearly shows that although there is no true condensate in the field, the field’s behaviour in its central region is dramatically different from that at larger radii. The presence of the frequency peak is indicative of the temporal coherence of this central region, and we will use the local autocorrelation function

$$G(\mathbf{x}, \tau; t_0) = [\mathfrak{F}_{t_0}^T]^{-1} \{H(\mathbf{x}, \omega; t_0)\}, \quad (9.19)$$

⁹It is important to note that such an artifact must arise in any field theory which is truncated either explicitly or implicitly, and that the defect occurring here is thus close to the minimal defect attainable in such a description.

of the field at point \mathbf{x} , at lag τ relative to the time t_0 to quantify the local temporal coherence of the field. We evaluate the local autocorrelation function by using the Wiener-Khinchin theorem [215].

In practice we take the discrete Fourier transform of $\tilde{H}(r_k, \omega; t_0)$ to obtain the azimuthally averaged autocorrelation $\tilde{G}(r_k, \tau; t_0)$, and we normalise it so that $|\tilde{G}(r_i, \tau = 0; t_0)| = 1$ for all radii r_i . We will use $|\tilde{G}(r_i, \tau; t_0)|$ to characterise the temporal coherence of the classical field, as a function of position, and take the time scale for its decay to represent the time scale over which the field remains temporally coherent at radius r_k , near time t_0 . Illustrative results are shown in figures 9.7(c) and (d).

At the centre of the trap (cyan (light grey) line in figure 9.7(d)), the peak of the autocorrelation function is broad and decays monotonically with increasing $|\tau|$. At larger radii ($2r_0 \lesssim r \lesssim 5r_0$), this peak possesses a nontrivial structure (black line), with $|\tilde{G}|$ vanishing and then increasing again as τ varies. Analysing the trajectories of vortices during the sampling period we conclude that this vanishing of the autocorrelation is due to the passage of vortices through these radii¹⁰. We identify the FWHM of the central peak (figure 9.7(d)) at a radius r_i as (twice) the correlation time of the classical field at this radius. At the trap centre, the correlation time $t_c \approx 2.2$ cyc, indicating that despite the presence of highly energetic thermal excitations and mode mixing in the central region, the temporal coherence of this region is significant. We note also that a small ridge appears in $|\tilde{G}(\tau)|$ for τ close to zero, and interpret this as representing the correlations of the thermal component of the field, and note that an analogous behaviour of *spatial* correlations in classical-field simulations was observed in [24].

The width of the broad envelope (figure 9.7(d)) is roughly constant until $r \approx 8r_0$, then decreases with increasing radius, and the strong peak in the power spectrum also rapidly reduces with radius for $r \gtrsim 8r_0$. Past this point the peak becomes much narrower (red (dark grey) line in figure 9.7(d)); the short correlation time of the outer cloud reflecting the broad range of energies characteristic of the thermally occupied excitations. Beyond the classical turning point, the evanescent nature of the basis modes yields a spurious amount of temporal coherence and this manifests as an increasing correlation time past this point (a similar phenomenon was observed for spatial correlations in [24]).

While the interpretation of the correlation time we calculate in this manner is complicated by the non-monotonic decay of $|\tilde{G}(\tau)|$ with increasing $|\tau|$, it nevertheless allows us to distinguish two qualitatively different regimes of behaviour in the classical-field equilibrium. In the outer region of the field, the short time scales over which correlations persist in the field indicate that this material constitutes a truly ‘thermal’ or ‘normal’

¹⁰We expect that averaging the autocorrelation over an ensemble of trajectories we would obtain a smoothly decaying function, due to the random distribution of vortex configurations in the ensemble.

fluid. In contrast, the central region of the field is spatially disordered, yet maintains temporal coherence over a comparatively long time scale. We identify the phase of the classical field in this central region as a *vortex liquid* [93, 152, 110, 121]: although the vortices are themselves disordered in their arrangement and motion, each supports a conserved supercurrent (phase winding of 2π). These supercurrents are robust in this region, in contrast to the normal phase of the field, where the spontaneous ‘pair production’ of vortices of opposite sign leads to the screening of the vortex interactions and ultimately the destruction of the conserved superflows [93, 9].

Away from equilibrium, the characterisation of the phases of the classical field is more subtle. Nevertheless, we can use temporal correlations to distinguish between different regimes of behaviour in the nonequilibrium field. For example, in figures 9.8(a) and (b) we plot the phase profile of the classical field at time $t = 100$ cyc and the corresponding correlation time, respectively. The correlation time maintains a near-constant value of $t_c \approx 1.1$ cyc as r increases until $r \approx r_0$, where the first vortex (phase singularity in figure 9.8(a)) occurs. At this point the correlation time begins to decay steeply with increasing r . Throughout the region $r < r_{tp}$ the phase structure is complicated, however we see that the correlation time of the central region is at least of order $t_c \sim 0.5$ trap cycles, indicating that despite the turbulent nature of the field here, interactions serve to maintain its coherence, and we identify the behaviour here as superfluid turbulence [19], in contrast to the outer region ($r \gtrsim 6r_0$), which has a very short correlation time ($t_c \lesssim 0.1$ cyc) and thus appears to be truly thermal material. This is in contrast to the interpretation presented by the authors of [194, 195], who considered the turbulence developed during the stirring process to be purely superfluidic (i.e., zero-temperature) in nature.

Relaxation of the vortex liquid

The characterisation of the classical field by its temporal correlations also allows us to quantify the relaxation of the central vortex-liquid material due to its interaction with the thermal cloud. In order to follow the relaxation of the vortex liquid, we measure the frequency ω_p of maximum occupation in the power spectra, for power spectra evaluated over various periods. Away from equilibrium the interpretation of this frequency is not entirely obvious, nevertheless it provides a measure of the energy of the quasi-coherent central structure, and we interpret it as representing an effective (quasi-)condensate chemical potential $\mu_p = \hbar\omega_p$ ¹¹. In figure 9.9 we plot this effective chemical potential of the

¹¹At the level of approximation of these arguments, we neglect any correction due to the finite population of the ‘condensate’ (cf. equation (5.14)).




Figure 9.8: Plots of (a) the condensate-band phase and (b) the corresponding correlation time as a function of radius, near time $t = 100$ cyc.




Figure 9.9: Effective chemical potential of the vortex liquid and thermal cloud. The data displayed is obtained from 10 trap cycle long samples spaced at 100 trap cycle intervals. Inset: Data with samples spaced at 10 trap cycle intervals over the initial decay (rise) of μ_p (μ_{th}).

vortex liquid, and also the chemical potential of the thermal cloud obtained using our fitting procedure (section 9.3.1) for comparison. We see that the vortex-liquid chemical potential reduces with time, as the liquid is damped by its coupling to the thermal cloud. By $t \approx 4000$ cyc, the two chemical potentials appear to have reached diffusive equilibrium, to the accuracy of our fitting procedure for μ and T .

9.3.4 Vortex motion

As noted in section 9.2, the motion of vortices (as viewed in the rotating frame) is initially very rapid, and slows as time goes on. In order to quantify this motion and its slowing with time, we track vortex trajectories, monitoring the coordinates of all vortices within a circular region of radius equal to the Thomas-Fermi radius of the initial lab-frame condensate. Due to the rotational dilation of the vortical fluid, this region always lies within the central vortex-liquid region, as defined by the temporal analysis described in section 9.3.3. In general, vortices enter and leave this counting region as time progresses. We are interested in the motion of vortices which persist in the counting region and can therefore be considered to exist in the condensate, rather than those which occur in the violently evolving condensate periphery. We therefore discard the trajectories of all

vortices which do not remain in the region for at least half the counting period of $T_v = 10$ trap cycles about time t_i .

The motion of the vortices which remain in the central region long enough to be considered is erratic on a length scale of order of the healing length, which is the origin of the thermal filling of the vortex core discussed in section 6.3.2. As this motion occurs on a short time scale, we interpret it as the signature of high-energy excitations of the vortices. Our interest however is in the low-energy component of vortex motion, which undergoes damping as the field relaxes to equilibrium. We therefore apply a Gaussian frequency filter centred on $\omega = 0$ to the trajectory components $\{x(t), y(t)\}$ to remove their highest *temporal* frequency components. We choose the width of this filter to be $\sigma_\omega \sim 0.6\omega_r$, and so the frequency components preserved in the trajectories correspond to low-energy excitations ($\omega \sim \omega_r$). From the N_t coordinate pairs (x_k, y_k) of the j^{th} filtered vortex trajectory we extract the mean speed along the trajectory arc,

$$v_j(t_i) = \frac{1}{T_v} \int ds \approx \frac{1}{T_v} \sum_{i=1}^{N_t-1} \sqrt{(x_i - x_{i-1})^2 + (y_i - y_{i-1})^2}. \quad (9.20)$$

We then average this quantity over the vortices included in the count, to find the mean speed per vortex $\bar{v}(t_i) = \sum_{j=1}^{N_v} v_j(t_i)/N_v$. In figure 9.10 we plot this quantity as a function of time. We note first that the mean vortex speed decays dramatically during the initial measuring period $t \approx 70 - 150$ cyc, as vortices enter the condensate with high velocities (with a substantial component opposite to trap rotation, as observed in the rotating frame), and their motion is damped heavily by their interaction with one another and with the thermal component. Subsequent to this, a secondary, slower damping is observed over the period $t \approx 150 - 1100$ cyc, after which the vortex motion seems to have reached a lower limit set by thermal excitations. During this period the condensate contains nearly the number of vortices it contains in equilibrium, and we interpret the damping during this period as the damping of low-energy excitations of the vortex liquid due to their interaction with thermally occupied, high energy modes. Taking the log of the speed data (inset to figure 9.10), we see that the decay during this period is approximately exponential, and performing a linear fit to the data over this period we extract a damping rate $\gamma_v \approx 6.3 \times 10^{-4} \text{ (cyc)}^{-1}$.

We note that the final distribution of vortices is disordered, with the vortices failing to become localised in a regular lattice. This result is consistent both with the considerations of section 9.3.1 and with pure-GPE studies of condensate stirring in 2D performed by Feder *et al.* [80], and Lundh *et al.* [171]. The results presented here (and those of [170]) suggest that the turbulent states observed in references [80, 171] are in fact finite-temperature




Figure 9.10: Mean vortex speed measured from the classical-field trajectory, revealing the rapid slowing of vortices during the initial stage of nucleation ($0 \lesssim t \lesssim 120$ cyc), and more gradual damping of the remaining collective excitations ($100 \lesssim t \lesssim 1000$ cyc). Inset: Logarithm of the mean vortex velocity (scaled by the minimum velocity measured \bar{v}_{\min}), and linear least-squares fit to the data over the damping period.

ones beyond the validity of the GP model employed in those calculations. By contrast, the calculations of Parker and Adams [194, 195] produced fully crystallised vortex lattices in 2D, indicating that additional damping mechanisms were available to the system in their simulations. It is possible that these are attributable to the spatial differencing technique [192] used in those calculations, however the precise origin of this error is not clear¹².

9.4 Dependence on simulation parameters

We compare now the behaviour of stirred condensates for different condensate sizes, i.e., for different values of the initial chemical potential μ_i . The field density and thus the nonlinear effects of mean-field interaction increase as the chemical potential is increased, and larger ‘condensates’ in rotational equilibrium will also contain larger numbers of vortices. In section 9.4.1 we assess the effect of varying the initial chemical potential on our classical-field solutions.

¹²While additional numerical procedures (amounting to evaporative cooling) were used in some calculations presented in references [195, 194], those additions were not essential for lattice crystallisation; i.e., relaxation to a lattice state was also observed in simulations performed without such additions [193].

It is also important in performing classical-field calculations such as those presented here to consider the effect of the cutoff height on the results of simulations. In section 9.4.2 we compare results for simulations with fixed initial chemical potential but varying cutoff heights, and quantify the effect of this purely technical parameter on the physical predictions of our model.

9.4.1 Chemical potential

The initial chemical potential impacts upon the system evolution both qualitatively, affecting the density profile of the condensate mode and thus the nature of the dynamical instability leading to vortex nucleation, and quantitatively, determining the equilibrium parameters of the stationary state of the classical field. The characteristics of the vortex array and the slowing dynamics of the vortices are also strongly dependent on the chemical potential. We consider these four aspects in turn.

Effect on dynamical instability

The chemical potential determines the strength of the mean-field repulsion influencing the condensate's shape. As the chemical potential becomes large, the condensate mode approaches that described by the Thomas-Fermi approximation, whereas for smaller chemical potentials the kinetic energy of the condensate becomes important, and the condensate boundary broadens. Consequently the visual distinction between the condensate mode and its incoherent excitations becomes less clear as the chemical potential is reduced. This affects the early evolution of the system and the onset of the dynamical instability in a pronounced manner. For the largest chemical potentials considered ($\mu_i \approx 17 - 20\hbar\omega_r$), the collective (quadrupole) excitations and their breakdown are clearly visible against a background of incoherent noise resulting from the initial vacuum occupation. For smaller chemical potentials, although quadrupole oscillations are still visible and the outer cloud becomes visibly more dense as time proceeds, indicating that material is indeed lost from the condensate, any ejection of material from the condensate is obscured by the blurred condensate boundary. Persistent 'ghost' vortices can be seen forming in the incoherent region about the condensate periphery very early in the evolution ($t \approx 20$ cyc for $\mu_i = 4\hbar\omega_r$). In figure 9.11 we plot the evolution of the field angular momentum per atom for three different initial chemical potentials¹³. The most obvious feature is that larger condensates support more angular momentum per atom, due to increased mean-field ef-

¹³Note that in each case the initial angular momentum of the field is nonzero, due to the contribution of the rotating-frame vacuum. However, the angular momentum of the vacuum is quickly dominated by that of the nascent thermal material.




Figure 9.11: Evolution of the field angular momentum for various initial chemical potentials.

fects. The evolution of the field angular momentum for the smallest condensate considered ($\mu_i = 4\hbar\omega_r$) reveals the onset of ‘irreversible’ (see section 3.3) coupling of angular momentum into the field at $t \approx 130$ cyc. Inspection of the field-density evolution shows that vortices begin to be nucleated into the central region of the field at about this time. We therefore conclude that despite the absence of visible break-up of the condensate, the unstable quadrupole oscillations still serve to produce the thermal component required to allow vortices to enter the central (initially) condensed region¹⁴.

Thermodynamic parameters

As described in section 9.3.1, a simple analysis based on energetic considerations allows us to predict the equilibrium temperatures and chemical potentials of our classical-field solutions, and the dependence of these parameters on the initial chemical potential and the condensate-band multiplicity of our simulations. In figure 9.12 we plot these analytic predictions alongside the values obtained using the fitting procedure of section 9.3.1. The dotted line shows the estimated chemical potential of the (idealised) norm-conserving,

¹⁴More generally, outside the dynamically unstable regime (e.g., at lower drive frequencies) the thermalisation of the added Wigner noise can eventually lead to the nucleation of vortices due to the *thermodynamic* instability of the condensate. This pathway to vortex formation is however much slower than that due to the dynamical instability, as the system must fluctuate over the energetic barrier to vortex nucleation (see section 2.3.3). In fact, in the absence of any trap anisotropy the angular momentum of the noise can be sufficient for the condensate to migrate to a state containing a vortex (cf. chapter 6).




Figure 9.12: Equilibrium temperatures and chemical potentials reached as a function of initial chemical potential, as determined by the fitting procedure. Dotted and dash-dotted lines indicate the predictions of equation (9.6) and equation (9.9) respectively. Numbers indicate times at which the measurements were made (thousands of trap cycles) in cases where numerical results were constrained by time. Points without numbers were measured at 10^4 cyc.

rotating-frame condensate mode, $\mu_\Omega = \mu_i \sqrt{1 - \Omega^2/\omega_r^2}$, while the dash-dotted line indicates the estimated equilibrium temperature given by equation (9.9). We note that despite its approximate (TF-limit) and asymptotic (zero thermal fraction) nature, the analytical estimate for μ is a reasonable prediction of the measured values. The measured value of μ is slightly higher than the TF prediction, and this discrepancy appears to increase with increasing μ_i , while the measured temperature appears consistently smaller than the analytical prediction at small μ_i , and exceeds it as the initial chemical potential is increased. However, the agreement seems very reasonable given the simplicity of the arguments presented in section 9.3.1, which neglect both the kinetic energy of vortex cores (TF approximation) [199], the depletion of the condensate population, and all other effects beyond a linear (Bogoliubov) description. We therefore conclude that our simple analysis captures the essential physics determining the thermodynamic parameters of the equilibrium. That is to say, the temperature is determined by the necessity of the system to redistribute the excess energy of the initial state (as viewed in the rotating frame), and that no additional heating of the system is required for the system to migrate to a state containing vortices.

Vortex-array structure

Perhaps the most striking difference between vortex arrays in simulations of differing condensate populations is the number of vortices present. As noted in section 9.3.2, the density of vortices is largely determined by the rotation rate of the condensate and is given approximately by equation (9.12), and (within the TF approximation) the total number of vortices in an equilibrium condensate is approximately proportional to the chemical potential μ .

The stability of a vortex array in equilibrium is determined by the competition between hydrodynamic and energetic considerations [75] (which favour the crystallisation of a rigid lattice), and the perturbing effect of thermal fluctuations in the atomic field. Close to zero temperature these fluctuations are interpreted as the thermal occupation of Bogoliubov excitations of the underlying vortex-lattice state. As the temperature of a vortex-lattice state is increased it eventually undergoes a transition to a disordered *vortex-liquid* state [93, 242, 110, 37] as the long-range order of the lattice is degraded by thermal fluctuations.

After a long time period, $t \sim 10^4$ trap cycles, all our simulated fields appear to describe states on the disordered side of this transition. All our solutions have vortices migrating within the central bulk, and leaving and entering through its periphery. In the smallest condensates, with the smallest vortex counts (i.e. $\mu_i = (4, 5, 6)\hbar\omega_r$ with counts $N_v \approx (3, 4, 5)$), the vortex array is small enough that *quasi-regular* configurations of vortices may occur, in which the vortex positions appear to fluctuate about approximate equilibrium positions on a triangular (or square) lattice. Such configurations may persist for as long as ~ 10 trap cycles, but ultimately break down and give way to new configurations as vortices cycle in and out of the central, high density region of the field. We interpret this cycling behaviour as the classical field ‘sampling’ different configurations during its ergodic evolution¹⁵.

Larger vortex arrays appear to be prohibited from forming regular structures over scales larger than the nearest-neighbour inter-vortex separations (the scale of the apparent order in the smaller condensates). As noted in section 9.3.1, the equilibrium temperature attained by our simulated atomic fields increases linearly with the initial chemical potential, and so it seems unlikely that increasing tendency to disordered behaviour would abate when the condensate size is increased further. In figures 9.13(a-c) we plot the coordinate-space densities of vortex arrays of differing sizes, corresponding to differing initial chemical potentials. Because of their small size compared to the extent of the condensate, vortex cores are difficult to experimentally image *in situ*, and so their

¹⁵An alternate (and contrary) explanation, that the field is exhibiting quasi-recurrent behaviour in the manner of [84], seems unlikely given that the trajectories appear to have converged to the temperatures predicted by the simple analysis presented in section 9.3.1.

presence is typically detected after free expansion of the atom cloud [173] by absorptive imaging techniques. For condensates in the TF regime the expansion is well described by a simple scaling of the position-space density [144, 48] (though the vortex core radius grows in somewhat greater proportion [173]). The Beer-Lambert law [222] for the absorption of an optical probe, and the assumption of a ground-state (Gaussian) profile in z (see section 4.1.1) yields for the transmitted intensity of the probe

$$I(x, y) = I_0 e^{-\gamma |\psi(x, y)|^2}, \quad (9.21)$$

with I_0 the input probe intensity and γ a constant which depends, in general, on the absorptance of the atomic medium. In figures 9.13(d-f) we plot simulated intensity profiles obtained using equation (9.21), where we have in each case simply set $\gamma = 2/\{|\psi|^2\}_{\max}$, with $\{|\psi|^2\}_{\max}$ the peak value of the density occurring in the distribution, so as to best represent the significant features. This process accentuates both the disparity in optical depth between the condensate and the outer thermal cloud, and the density fluctuations in the central condensate region, in contrast to the logarithmic plots (figures 9.13(a-c)) which suppress them. The images show the thermally distorted vortex-array structure that might be measured in an experiment. The fact that the formation of rigid vortex lattices here is inhibited by thermal fluctuations, while ordered lattices have been formed experimentally at comparatively high temperatures [175, 1], suggests that this inhibition of lattice order may be a manifestation of the well-known susceptibility of two-dimensional systems to long-wavelength fluctuations [214]. Indeed two theoretical studies [208, 242] of vortex-lattice melting suggest that the melting temperature is very low in quasi-2D trapping geometries¹⁶.

Damping of vortex motion

While the vortex motion as measured by the mean velocity defined in section 9.3.4 exhibits the same overall damping behaviour in all our simulations, the damping appears to be generally nonexponential and quite sensitive to the initial noise conditions. We have observed additional features (e.g. transient and oscillatory behaviours superposed with the damping) peculiar to each trajectory. It is therefore difficult to unambiguously quantify the rates of motional damping in order to compare the dependence on initial chemical potential, without generating ensemble data for each parameter set, which would be a very heavy task numerically. We therefore plot in figure 9.14(a) only the behaviour of three particular trajectories, which indicate at a qualitative level that the vortex motion

¹⁶The group of Cornell have performed experiments on the vortex-lattice melting transition in trapped Bose condensates [155], however the results of these experiments appear to remain unpublished.




Figure 9.13: (a-c) Coordinate-space densities of the classical fields with initial chemical potentials $\mu_i = 5, 10, 18\hbar\omega_r$ at times $t = 9900, 9900, 8800$ cyc respectively. (d-f) Simulated absorption images generated from the same densities.

in the smaller condensates damps to equilibrium more quickly. We note however that the level of vortex motion that each trajectory damps to shows a clear dependence on the initial chemical potential. In figure 9.14(b) we plot the mean velocity of vortices over the period $t = 1900\text{--}2000$ cyc, at which point the vortex motion in each simulation appears to have essentially relaxed to its equilibrium level. The mean velocity so measured clearly increases with chemical potential. After comparing the results for varying cutoff at fixed initial chemical potential (not shown) we conclude that this is likely due to the increase in equilibrium temperature with increasing initial condensate size.

9.4.2 Cutoff dependence

Of crucial importance in classical-field simulations is the effect of the cutoff energy (and thus the basis size) upon the system behaviour. The dependence of the simulation results on the cutoff can only be eliminated by including above-cutoff physics, and no prescription for doing so in general nonequilibrium scenarios is currently known. However, the cutoff ‘height’ is chosen on physical grounds (so that all modes significantly modified by interactions are included in the dynamics), and so while we expect (e.g.) the equilibrium temperature to depend on the precise cutoff chosen (see equation (9.9)), we do not expect the dynamics to differ qualitatively between simulations with different cutoff heights. It might be expected [170], however, that the cutoff dependence of the temperature will influence relaxation rates, which depend strongly on the thermal occupation of system excitations (see, e.g., [82]).

Thermodynamic parameters

In figure 9.15 we plot the temperatures and chemical potentials attained in simulations with $\mu_i = 20\hbar\omega_r$ and cutoff heights E_R varying over the range $2\mu_i \leq E_R \leq 3\mu_i$. We note first that while the temperature shows a clear dependence on E_R , the chemical potential appears relatively insensitive to it. The system therefore reaches a chemical potential closest to that resulting from the idealised transformation of the condensate to a lattice state without loss of atoms from the condensate mode. The measured temperature decreases with increasing cutoff as expected, but differs from the $T \propto 1/(\mathcal{M}-1)$ behaviour predicted for a classical field in the Bogoliubov limit. We attribute this to the constraint imposed by the low cutoff, which artificially constrains the natural modes of the system, leading to spurious thermodynamic behaviour. Indeed for $E_R = 2\mu$, the classical turning point $r_{\text{tp}} \sim 1.7R_{\text{TF}}$, where $R_{\text{TF}} \approx 7.8r_0$ is the TF radius of the lattice state, and so the corruption the quasiparticle mode shapes may be quite pronounced.

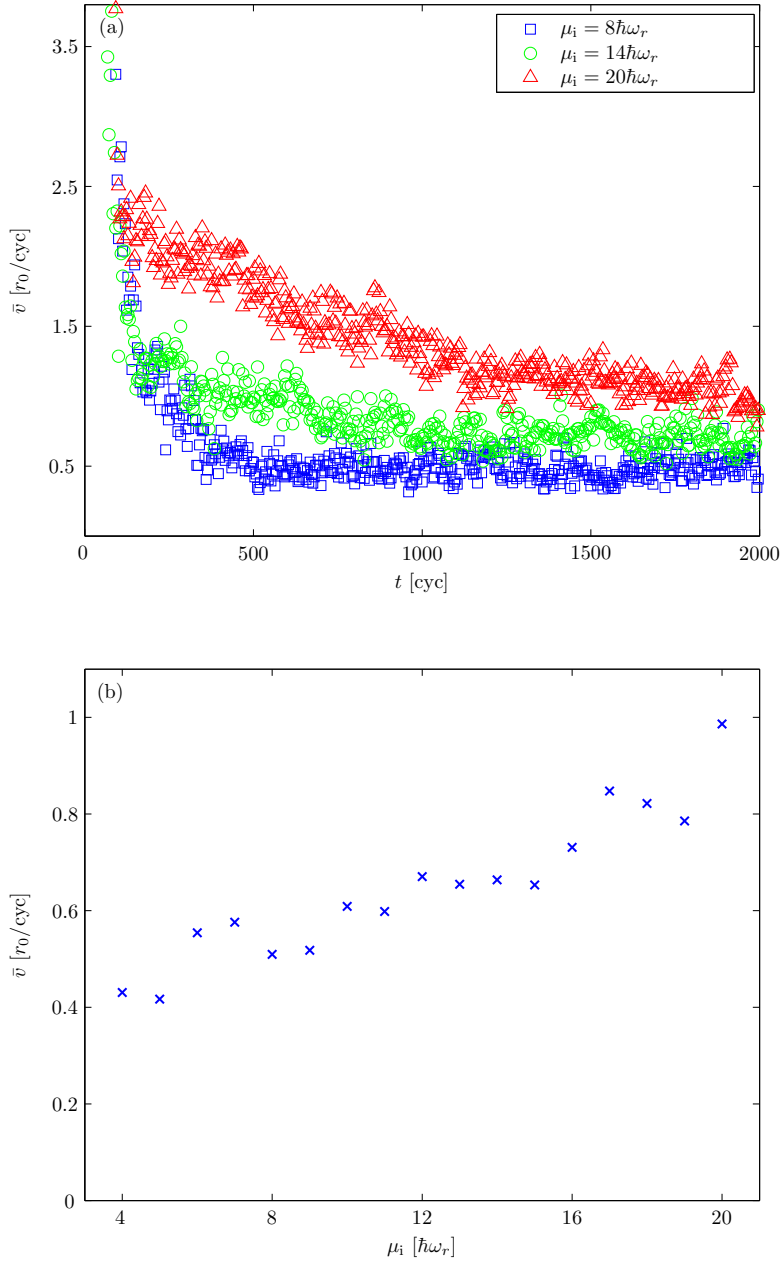


Figure 9.14: (a) Vortex motional damping for three representative choices of initial chemical potential μ_i . (b) Mean vortex speed (averaged over the period $t = 1900\text{--}2000$ cyc) as a function of initial chemical potential μ_i .)




Figure 9.15: Field temperatures and chemical potentials reached as a function of cutoff. Dashed and dash-dotted lines indicate the predictions of equation (9.6) and equation (9.9), respectively. Measurement times are indicated as in figure 9.12.

Damping

As noted in section 9.3.4, it is difficult to extract quantitative comparisons of relaxation rates from the damping of vortex motion. Here we consider the relaxation of another quantity: the effective chemical potential μ_p (power spectrum peak) of the vortex liquid, introduced in section 9.3.3. In figure 9.16, we show the decay of this quantity with time for three values of the cutoff energy E_R . We observe that the damping is reasonably insensitive to the cutoff. A possible explanation for this behaviour is that as the effect of lowering the cutoff is to decrease the number of thermally occupied modes in the system and consequently to increase the temperature of these modes, that these two effects serve to approximately cancel one another in determining the rates of thermal damping. An alternate possibility is that the relaxation dynamics of the low energy modes in the system are relatively insensitive to the precise parameters of the high-energy modes due to the low dimensionality of the system [209]. It is interesting to note, however, that the results of the Ketterle group on the formation of vortices in three dimensions [1] suggest that the rate of relaxation to the lattice state is insensitive to the temperature of the cloud. Determining the mechanisms underlying this somewhat counterintuitive result would require an analysis of fully three-dimensional classical-field simulations, and we leave this heavy computational task for future work.




Figure 9.16: Decay of the effective chemical potential of the vortex liquid (as defined in section 9.3.3) for trajectories with initial chemical potential $\mu_i = 20\hbar\omega_r$ and three different values of the energy cutoff E_R .

9.5 Summary

We have carried out the first strictly Hamiltonian simulations of vortex nucleation in stirred Bose-Einstein condensates. Our approach is free from grid-method artifacts such as aliasing and spurious damping at high momenta, and carried out to a very high and well controlled numerical accuracy, enabling a controlled study of field thermalisation and vortex nucleation within an explicitly Hamiltonian classical-field theory.

Vacuum symmetry breaking.—The importance of symmetry breaking has been highlighted in previous works [200, 194, 195]. Sampling of initial vacuum noise provides an irreducible mechanism for breaking the twofold rotational symmetry of the classical-field solution.

Thermalisation.—Resonant excitation of the quadrupole instability evolves the system from a zero-temperature initial state through a dynamical thermalisation phase in which a rotating thermal cloud forms. Subsequent vortex nucleation in the remaining condensate shows that the cloud provides the requisite dissipation to evolve the condensate subsystem toward a new quasi-equilibrium state in the rotating frame. Although the classical-field evolution is Hamiltonian, the stochastically evolving high-energy modes of the classical field provide the thermal-cloud damping that previous authors [254, 200] had identified as necessary to drive the nucleation of vortices and the relaxation of the vortex array.

Identifying the thermal cloud.—We have quantified the heating caused by the dynamical instability, using a range of measures. Spectral analysis of the nonequilibrium classical-field evolution has been used to identify the condensate chemical potential, and thermal cloud properties were extracted using a semiclassical fit at high energies. While the chemical potentials of the condensate and emerging thermal cloud are slower to equilibrate, the moments of inertia, temperature, and angular momentum of the classical field are rapidly transformed by the instability to those of a rotating, heated Bose gas.

Frustrated crystallisation.—Vortices are seen to nucleate and enter the bulk of the condensate but a regular Abrikosov lattice does not form. Instead the vortices distribute throughout the condensate in spatially disordered vortex-liquid state, consistent with substantial thermal excitation of an underlying regular vortex lattice state. In this respect our results are consistent with previous 2D GPE treatments of similar stirring configurations, which did not observe vortex-lattice crystallisation [80, 171].

We conclude that vortex nucleation by stirring is a finite-temperature effect, which arises from dynamical thermalisation when initiated from a zero-temperature Bose condensate. While we observe vortex *nucleation* in 2D, the asymptotic absence of lattice rigidity suggests that the temperature attained in accommodating the excess (rotating-frame) energy of the initial state is such that vortex lattice *crystallisation* is inhibited in this conserving stirring scenario in 2D.

Chapter 10

Conclusion

10.1 Thesis summary

In this thesis, we have conducted extensive studies using the *classical-field* method for the description of equilibrium and nonequilibrium finite-temperature degenerate Bose-gas systems. We have developed new approaches to interpreting the results of classical-field simulations in equilibrium and nonequilibrium regimes, and we have considered the finite-temperature dynamics of quantum vortices, their decay and the dynamics of their formation in the classical-field formalism.

We examined the temporal behaviour of equilibrium classical-field trajectories, and investigated the relationship between temporal correlations and other measures of coherence in the field. We demonstrated that on short time scales, the condensate in the field can be identified as a nonzero *first moment* of the field in an appropriate phase-rotating frame. We thus showed how an analogue of the symmetry-breaking assumption typically employed in mean-field theories of Bose condensation emerges naturally from a classical-field trajectory. We found that the effective eigenvalue of the condensate mode so obtained exhibits the expected temperature-dependent relation to the thermodynamic chemical potential of the field. We also demonstrated the generality of the symmetry-breaking interpretation, by calculating the *anomalous* thermal-field correlations comprising the pair matrix, and explicitly constructing the anomalous thermal density of the field, which we found to have spatial form and temperature dependence in good agreement with the results of mean-field theories.

Turning our attention to the dynamics of quantum vortices in finite-temperature classical-field theory, we showed that equilibrium precessing-vortex configurations of the classical field could be obtained from the ergodic evolution of classical-field trajectories constrained to a finite conserved angular momentum. These simulations each exhibited a

quantum vortex precessing in rotational equilibrium with the rotating thermal component of the field, naturally realising the balance of coherent and frictional forces on the vortex. We found that as an off-axis vortex state breaks the rotational symmetry of the Hamiltonian, the vortical condensate mode is fragmented over angular momentum eigenmodes in the formal microcanonical correlations of the field. We demonstrated by abandoning formal microcanonical (ergodic) averages, and instead considering the field fluctuations on shorter time scales, that the symmetry-broken condensate orbital is revealed by its coherent, number-fluctuation-suppressed statistics. In addition, this approach revealed the thermal core-filling of the vortex, and produced also the spurious (Goldstone) mode corresponding to the broken rotational symmetry, which we found to have anomalous fluctuation statistics appropriate to its spurious nature. We showed that the angular momentum of the condensed mode decreases dramatically with increasing field energy (and thus temperature) at fixed overall angular momentum of the field, and that the vortex precessional radius and frequency correspondingly increase until the vortex leaves the condensate. We demonstrated that the decomposition of the classical field into condensed and noncondensed parts on the basis of our short-time analysis allowed us to calculate the rotational properties of each component, and we found qualitative agreement between the angular velocity of the thermal cloud and the precessional frequency of the vortex.

We then applied this methodology to a nonequilibrium scenario of Bose-field dynamics: the arrest of a finite-temperature precessing-vortex configuration by a fixed trap anisotropy. Our classical-field approach provided the first simulations of vortex arrest which include the coupled nonequilibrium dynamics of the condensate and thermal cloud, describing the migration of an *equilibrium* rotating finite-temperature field configuration to a new, irrotational equilibrium, due to the action of an externally applied potential. Decomposing the nonequilibrium field into condensed and noncondensed components, we observed the dynamics of angular momentum exchange between the condensate and thermal cloud, and the loss of angular momentum from the field due to the trap anisotropy. We found that the anomalous rotational response of the single-vortex state leads to a counter-intuitive spin-up of *both* the condensate and the thermal component during the decay, and we observed nonequilibrium oscillations in the transfer of angular momentum between the two components. We quantified the heating of the atomic field during its relaxation, and found it to be commensurate with the conversion of rotational kinetic energy into thermal energy during the arrest of the field's rotation. By varying the severity of the trap anisotropy, we altered the relative strength of the cloud-trap coupling, and observed qualitatively different regimes of decay. We found that for weak trap anisotropies, the thermal field persists in a long-lived nonequilibrium steady state, with rotation rate intermediate between that of the condensate and that of the static trap. We observed in

this regime the slow loss of angular momentum from the condensate mode while the angular momentum of the thermal field remained approximately constant, until the vortex neared the condensate boundary. For stronger trap anisotropies, we observed strongly nonequilibrium dynamics, in which the angular momentum of the thermal cloud may be almost entirely depleted before the condensate responds. Comparing the relaxation of the thermal field following the loss of the vortex from the condensate across simulations with different ellipticities, we found the scaling of the cloud spin-down time with ellipticity to be in good agreement with the predictions of a Boltzmann-gas model, demonstrating that our classical-field treatment correctly describes the dynamics of the thermal material.

Finally, we considered the problem of the stirring of condensates to produce vortex lattices. While there has been extensive theoretical work to describe the condensate-stirring experiments, there remain many problems with previous treatments. We therefore considered the problem within our classical-field approach, in order to obtain a more comprehensive and unified treatment. We simulated the stirring of quasi-two-dimensional condensates that were initially at zero temperature, and we included a representation of vacuum fluctuations in the initial conditions. We showed that material spontaneously ejected from the condensate bulk during its dynamic instability quickly comes to a local rotational and thermal equilibrium, forming a rotating thermal cloud which drives the growth of condensate surface modes, leading to the nucleation of vortices. We showed that the motion of the initially turbulent swarm of vortices admitted into the central bulk of the field in this process is slowly damped by the frictional effect of the thermal field. We characterised the evolution of the thermodynamic parameters of the field during this relaxation process, and showed that the equilibrium temperature of the field is well predicted by simple energy-conservation arguments. We found that the heating of the field is such that the vortices in this two-dimensional scenario are prohibited from crystallising into a rigid lattice by thermal fluctuations, and instead form a disordered *vortex liquid*. We showed that the disordered motion of vortices in this phase destroys true condensation in the equilibrium field, but demonstrated that the quasi-coherent vortex liquid could be distinguished from the remaining thermal component of the field by its *temporal* correlations.

10.2 Directions for future work

The approaches to characterising the coherence of the classical field by its temporal correlations we have developed in this thesis provide a useful extension of the classical-field method. It would be interesting to further develop this approach to obtain classical-field

analogues of temporal correlation functions, such as the one-body Green's functions, that appear in more traditional descriptions of Bose-condensation. It would also be helpful to derive formal signatures of (e.g.) quasicondensation and other more subtle coherence features in terms of the temporal correlations of classical-field trajectories. In our approach we have exploited the effective breaking of the phase symmetry by the condensate orbital to define an anomalous averaging operation. In principle, however, we expect all meaningful spatial (or modal) correlations of the equilibrium field to be able to be expressed in terms of *observable* combinations of classical-field variables; i.e., combinations which are invariant under global gauge rotations of the classical-field phase. It would therefore be of interest to furnish a characterisation of the anomalous classical-field fluctuations in such terms. It would also be interesting to see if the normal modes of the classical-field can be obtained from these correlations, and moreover to calculate (for example) the triplet correlations which are related to the scattering of particles in and out of the condensate mode, and the higher cumulants of the field which characterise the beyond-Gaussian corrections to the field fluctuations we might expect to find in the nonperturbative classical-field description.

The simulations of vortex dynamics in this thesis are restricted to quasi-2D systems, and we have considered only a strictly Hamiltonian description of the low-energy coherent region. The extension of the Gauss-Laguerre quadrature technique for rotating systems to a 3D method would allow us to consider the more complicated dynamics of vortex lines of finite extent. The extension to three dimensions is conceptually trivial due to the separable nature of the eigenstates of the single-particle Hamiltonian, however, the calculations would be significantly more demanding numerically. This extension would allow us to study the 3D dynamics of condensate stirring, and to ascertain the role of dimensionality in the vortex-lattice relaxation process. The inclusion of a mean-field description of above-cutoff atoms would extend the equilibrium rotating-frame PGPE method to allow quantitative comparisons with experiments. The inclusion of the *dynamics* of the above-cutoff region and its coupling to the coherent region in nonequilibrium regimes such as the vortex arrest and condensate stirring scenarios remains an outstanding goal for the classical-field methodology. It would also be interesting to further characterise the equilibrium vortex-liquid phase, identify the lattice-liquid and liquid-normal-fluid phase boundaries, and map out the (classical) phase diagram of the rotating Bose gas. Finally, it would be interesting to see what insights into nonequilibrium and steady-state turbulence in degenerate Bose-gas systems can be obtained from the family of classical-field methods.

Appendix A

Condensation in the noninteracting classical field

In this appendix we illustrate the analogy between the classical-field system and the Bose-field system by considering the thermodynamics of the classical-field distribution. Specifically, we consider the *noninteracting* limit of the classical field, and derive its transition temperature and the temperature dependence of the condensed fraction of the field.

Our approach here follows the discussion of the noninteracting Bose gas presented in section 2.2.1. We introduce the classical-field distribution (see section 3.3.2)

$$f^{\text{CF}}(\epsilon_\nu) = \frac{k_{\text{B}}T}{\epsilon_\nu - \mu} \Theta(E_R - \epsilon_\nu), \quad (\text{A.1})$$

where μ is the chemical potential, ϵ_ν is the energy of a particular single-particle level, E_R is the cutoff energy, and $\Theta(x)$ denotes a Heaviside function. As in the case of the full Bose distribution, the occupation of the ground level $f^{\text{CF}}(\epsilon_0 \equiv 0)$ becomes negative for $\mu > 0$, and so we require $\mu \leq 0$. Consequently the population of any excited state (below the cutoff) is bounded from above: $f^{\text{CF}}(\epsilon_\nu) \leq k_{\text{B}}T/\epsilon_\nu$. The critical temperature is thus that temperature at which the population of excited states is equal to the total classical-field population N , for $\mu = 0$. We follow [60, 204] by making the semiclassical approximation; i.e., we neglect the discrete level structure of the single-particle energy, and treat it as a continuous variable, introducing the general density of states $g(\epsilon) = C_\alpha \epsilon^{\alpha-1}$, and we find¹

$$N = \int_0^{E_R} d\epsilon g(\epsilon) \frac{k_{\text{B}}T_c}{\epsilon} = \frac{C_\alpha E_R^{\alpha-1} k_{\text{B}}T_c}{\alpha - 1}. \quad (\text{A.2})$$

¹For $\alpha > 1$, the condensate (with $\epsilon = 0$) does not contribute to the integral, and so we can safely take the lower bound of the integral to zero (see discussion in reference [197]).

Appendix A. Condensation in the noninteracting classical field

The critical temperature of the classical field is therefore

$$k_{\text{B}}T_{\text{c}} = \frac{\alpha - 1}{C_{\alpha}E_R^{\alpha-1}}N, \quad (\text{A.3})$$

which, of course, depends explicitly on the cutoff energy E_R . We note that, as in the case of Bose statistics, the integral diverges for $\alpha \leq 1$. Recalling that $\alpha = d/2 - 1$ for a d -dimensional homogeneous system and $\alpha = d - 1$ for a d -dimensional harmonic trap [203], the classical-field thermodynamics thus preserve the result of the full Bose system that no condensation occurs in a homogeneous system for $d \leq 2$, nor in a 1-dimensional harmonic trap, in the thermodynamic limit. Moreover, the excited-state population at a temperature $T < T_{\text{c}}$ is given by expression (A.2) with $T_{\text{c}} \rightarrow T$, and we thus find

$$\frac{N_0}{N} = 1 - \frac{T}{T_{\text{c}}}, \quad (\text{A.4})$$

as found by Connaughton *et al.* [57] for the homogeneous classical-wave system in three dimensions. We note, however, the somewhat surprising result that the relation (A.4) is independent of geometry and dimensionality, provided that the integral in (A.2) converges. In the limit that the distribution $f^{\text{CF}}(\epsilon_{\nu})$ coincides with the true Bose-Einstein distribution throughout the below-cutoff region, the corresponding *above-cutoff* population of course rapidly increases with temperature [64], and depends explicitly on the geometry, so that the correct geometry-dependent scaling is regained for the fraction $N_0/(N + N_{\text{above}})$.

Appendix B

Interaction-picture integration algorithms

In this appendix we review the *interaction-picture* methods for the temporal integration of Gross-Pitaevskii-like equations. These methods are all descendants of the original RK4IP algorithm developed by Ballagh [17], and the adaptive ARK45-IP variant of the algorithm due to Davis [63].

B.1 The RK4IP algorithm

The fundamental issue we face in evolving (nonlinear) Schrödinger-like equations is the efficient evaluation of operators which are nonlocal in their action on the wave function. A key observation is that the operators which are nonlocal in *position* space are typically local in *momentum* space; for example, the Laplacian $-\nabla^2$ becomes the multiplication $k^2 \times$ in momentum space. This observation forms the basis for Fourier-based integration methods, such as the symmetrised split-step method (see [141] and references therein), which alternates between the two representations of the wave function, applying operators which are local in each space. The original RK4IP algorithm [17] is in essence an efficient extension of the split-step method, in which the Fourier transform is used to transform to an *interaction picture* with respect to the Laplacian, and the effect of the remaining operators (which are usually local) is evaluated in this picture, using a standard Runge-Kutta approach. In general, *projected* Gross-Pitaevskii equations are formulated in a Galerkin approach, i.e., the field is propagated in the representation which diagonalises the Laplacian, or some more general base Hamiltonian H_0 (see section 4.1.2). In such cases the basis amplitudes are transformed to the interaction picture with respect to

Appendix B. Interaction-picture integration algorithms

H_0 , exploiting its diagonal form in this representation, and the effect of the remaining (projected) terms on the basis amplitudes is evaluated in this interaction picture.

B.1.1 Fourth-order Runge Kutta

Runge-Kutta (RK) methods form a standard approach [215] for the integration of a system of differential equations

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t), \quad (\text{B.1})$$

in steps of Δt . To advance the solution of equation (B.1) from $\mathbf{x}(t_i)$ to $\mathbf{x}(t_i + \Delta t)$, the fourth-order RK method calculates four estimates of the function change over the interval $[t_i, t_i + \Delta t]$

$$\mathbf{k}_1 = \Delta t \mathbf{f}(\mathbf{x}(t_i), t_i) \quad (\text{B.2a})$$

$$\mathbf{k}_2 = \Delta t \mathbf{f}(\mathbf{x}(t_i) + \mathbf{k}_1/2, t_i + \Delta t/2) \quad (\text{B.2b})$$

$$\mathbf{k}_3 = \Delta t \mathbf{f}(\mathbf{x}(t_i) + \mathbf{k}_2/2, t_i + \Delta t/2) \quad (\text{B.2c})$$

$$\mathbf{k}_4 = \Delta t \mathbf{f}(\mathbf{x}(t_i) + \mathbf{k}_3, t_i + \Delta t), \quad (\text{B.2d})$$

and averages these to produce the final estimate

$$\mathbf{x}(t_i + \Delta t) = \mathbf{x}(t_i) + [\mathbf{k}_1 + 2(\mathbf{k}_2 + \mathbf{k}_3) + \mathbf{k}_4]/6. \quad (\text{B.3})$$

The error term in this estimate is of order $(\Delta t)^5$, and the method is therefore fourth-order accurate.

B.1.2 RK4IP

The projected-GP equations we wish to solve are of the general form

$$i \frac{d\mathbf{c}}{dt} = H_0 \mathbf{c} + \hat{G}(t) \mathbf{c}, \quad (\text{B.4})$$

where the vector \mathbf{c} lists all the mode coefficients c_j , and the nonlinear operator $\hat{G}(t)$ includes the projected nonlinear interaction term and possibly some other perturbing term (see section 3.2.2). The integration relies on the transformation of the mode coefficients to an interaction picture with origin at some time t_0 , i.e.

$$\mathbf{c}^I(t) = e^{iH_0(t-t_0)} \mathbf{c}(t). \quad (\text{B.5})$$

The resulting differential equation in the interaction picture is

$$i \frac{d\mathbf{c}^I}{dt} = \hat{G}^I \mathbf{c}^I, \quad (\text{B.6})$$

where the interaction-picture nonlinear operator

$$\hat{G}^I(t) = e^{iH_0(t-t_0)} \hat{G} e^{-iH_0(t-t_0)}. \quad (\text{B.7})$$

For each step $\mathbf{c}(t_i) \rightarrow \mathbf{c}^I(t_i + \Delta t)$ of the integration, we choose a new interaction-picture origin t_0 , transform to the interaction picture, evaluate the four Runge-Kutta slopes \mathbf{k}_i , and transform back. A considerable simplification is obtained by choosing the interaction picture origin $t_0 = t_i + \Delta t/2$, so that the nonlinear operator need not be transformed to evaluate the midpoint slopes \mathbf{k}_2 and \mathbf{k}_3 ¹. The resulting algorithm is

$$\mathbf{c}^I(t_i) = e^{-i(\Delta/2)H_0} \mathbf{c}(t_i) \quad (\text{B.8a})$$

$$\mathbf{k}_1 = e^{-i(\Delta/2)H_0} [-i\Delta \hat{G}(t_i)] \mathbf{c}(t_i) \quad (\text{B.8b})$$

$$\mathbf{k}_2 = -i\Delta t \hat{G}(t_i + \Delta t/2) [\mathbf{c}^I(t_i) + \mathbf{k}_1/2] \quad (\text{B.8c})$$

$$\mathbf{k}_3 = -i\Delta t \hat{G}(t_i + \Delta t/2) [\mathbf{c}^I(t_i) + \mathbf{k}_2/2] \quad (\text{B.8d})$$

$$\mathbf{k}_4 = -i\Delta t \hat{G}(t_i + \Delta t) e^{-i(\Delta/2)H_0} [\mathbf{c}^I(t_i) + \mathbf{k}_3] \quad (\text{B.8e})$$

$$\mathbf{c}(t_i + \Delta t) = e^{-i(\Delta/2)H_0} \{ \mathbf{c}^I(t_i) + [\mathbf{k}_1 + 2(\mathbf{k}_2 + \mathbf{k}_3)]/6 \} + \mathbf{k}_4/6. \quad (\text{B.8f})$$

We note also that this fixed-step-size interaction-picture RK approach can be adapted to the integration of stochastic differential equations [101].

B.2 The ARK45-IP algorithm

One issue with the fixed-step-size RK4IP method presented in section B.1.2 is that of choosing the appropriate step size Δt . A step size which is too large leads to numerical inaccuracy and ultimately to the failure of the integration algorithm to produce the correct dynamics for the classical field. However, the necessity of choosing a sufficiently short time step must be contrasted with the dependence of the total simulation time on the step size. In general, one would like to take the largest step size which is still consistent with accurate integration of the field trajectory. Choosing such a step size is particularly

¹In the original RK4IP implementation the solution was propagated in real space, and so the interaction-picture transformation $e^{iH_0 t}$ was the most numerically demanding task. By contrast, the modern implementations propagate the solution in the spectral basis so that the application of \hat{G} is the numerically demanding operation.

difficult in nonequilibrium scenarios such as we consider in chapters 7 and 9, in which the composition of the field (e.g., the occupation of high-energy modes) and thus the appropriate step size may change during the evolution of the field. This issue can be solved by using an *adaptive-step-size* algorithm, which allows one to choose an accuracy tolerance at the beginning of the integration. Throughout the integration the adaptive algorithm then shortens or lengthens the step size so as to always take the largest step consistent with the prescribed accuracy requirement.

B.2.1 Adaptive Runge-Kutta

This algorithm is a generalisation of the RK method, based on the notion of calculating slopes of the function at trial midpoints that can be combined to give *two* estimates of the function value at the step endpoint $\mathbf{x}(t_i + \Delta t)$, of two different orders. By comparing these two estimates we thus obtain a measure of the truncation error in the higher-order estimate, and this allows us to choose after each step how to proceed with the integration; i.e., whether to accept the step, or reject the step and make a second attempt with a shortened step size. The most common adaptive RK algorithm is based on the fifth-order RK formula of form

$$\mathbf{k}_1 = \Delta t \mathbf{f}(\mathbf{x}(t_i), t_i) \quad (\text{B.9a})$$

$$\mathbf{k}_2 = \Delta t \mathbf{f}(\mathbf{x}(t_i) + b_{21}\mathbf{k}_1, t_i + a_2\Delta t) \quad (\text{B.9b})$$

...

$$\mathbf{k}_6 = \Delta t \mathbf{f}(\mathbf{x}(t_i) + b_{61}\mathbf{k}_1 + \dots + b_{65}\mathbf{k}_5, t_i + a_6\Delta t) \quad (\text{B.9c})$$

$$\mathbf{x}(t_i + \Delta t) = \mathbf{x}(t_i) + d_1\mathbf{k}_1 + d_2\mathbf{k}_2 + d_3\mathbf{k}_3 + d_4\mathbf{k}_4 + d_5\mathbf{k}_5 + d_6\mathbf{k}_6, \quad (\text{B.9d})$$

where the a_i , b_{ij} , and d_i are rational-number coefficients². The key point is that having evaluated the slopes \mathbf{k}_i , a fourth-order estimate

$$\tilde{\mathbf{x}}(t_i + \Delta t) = \mathbf{x}(t_i) + \tilde{d}_1\mathbf{k}_1 + \tilde{d}_2\mathbf{k}_2 + \tilde{d}_3\mathbf{k}_3 + \tilde{d}_4\mathbf{k}_4 + \tilde{d}_5\mathbf{k}_5 + \tilde{d}_6\mathbf{k}_6, \quad (\text{B.10})$$

can also be obtained. The difference between the two is then taken as an estimate of the truncation error in the fifth-order estimate $\mathbf{x}(t_i + \Delta t)$. Since $\mathbf{x}(t)$ is in general a vector, we choose some summary statistic such as (for example) $\Delta_m = \max_j \{|\tilde{x}_j(t_i + \Delta t) - x_j(t_i + \Delta t)|\}$ as an estimate of the error. With a given error tolerance Δ_{tol} in mind, we can use the expectation that Δ_m scales as h^5 to deduce the step size that should be taken. If the error

²The particular values of these coefficients that we use here are the Cash-Karp values, which are tabulated in reference [215].

is too large, the step size is reduced, while if it is smaller than the required tolerance, the step size can be safely expanded. We are thus lead to choose the step size according to

$$\Delta t_{n+1} = \begin{cases} S\Delta t_n \left| \frac{\Delta_{\text{tol}}}{\Delta_m} \right|^{1/5} & \Delta_m \leq \Delta_{\text{tol}} \\ S\Delta t_n \left| \frac{\Delta_{\text{tol}}}{\Delta_m} \right|^{1/4} & \Delta_m > \Delta_{\text{tol}}. \end{cases} \quad (\text{B.11})$$

The safety factor S here is a coefficient a few percent smaller than unity, and the rationale for the choice of exponents in equation (B.11) is given in reference [215].

B.2.2 ARK45-IP

The ARK45-IP algorithm is due to Davis [63], and provides an adaptive variant of the RK4IP algorithm, by applying the formula of equations (B.9a-B.11) in the interaction picture with respect to H_0 . The resulting algorithm is more complicated than the fixed-step RK4IP algorithm, due to the range of intermediate times $a_i\Delta t$ at which the RK slopes are evaluated. Thus no simplifications of the interaction-picture transformations can in general be made in this case. However due to the comparatively light numerical load of applying the operator $e^{i(a_i\Delta t)H_0}$ in the spectral representation (in which it is diagonal), this is of little concern. The algorithm for the fifth-order estimate is of the general form

$$\mathbf{k}_1 = -i\Delta t \hat{G}(t_i) \mathbf{c}(t_i) \quad (\text{B.12a})$$

$$\mathbf{k}_2 = e^{-i(a_2\Delta t)H_0} \left[-i\Delta t \hat{G}(t_i + a_2\Delta t) \right] e^{i(a_2\Delta t)H_0} [\mathbf{c}(t_i) + b_{21}\mathbf{k}_1] \quad (\text{B.12b})$$

...

$$\begin{aligned} \mathbf{k}_6 &= e^{-i(a_6\Delta t)H_0} \left[-i\Delta t \hat{G}(t_i + a_6\Delta t) \right] e^{i(a_6\Delta t)H_0} \\ &\quad \times [\mathbf{c}(t_i) + b_{61}\mathbf{k}_1 + b_{62}\mathbf{k}_2 + b_{63}\mathbf{k}_3 + b_{64}\mathbf{k}_4 + b_{65}\mathbf{k}_5] \end{aligned} \quad (\text{B.12c})$$

$$\mathbf{c}(t_i + \Delta t) = \mathbf{c}(t_i) + d_1\mathbf{k}_1 + d_2\mathbf{k}_2 + d_3\mathbf{k}_3 + d_4\mathbf{k}_4 + d_5\mathbf{k}_5 + d_6\mathbf{k}_6, \quad (\text{B.12d})$$

and the fourth-order estimate is

$$\tilde{\mathbf{c}}(t_i + \Delta t) = \mathbf{c}(t_i) + \tilde{d}_1\mathbf{k}_1 + \tilde{d}_2\mathbf{k}_2 + \tilde{d}_3\mathbf{k}_3 + \tilde{d}_4\mathbf{k}_4 + \tilde{d}_5\mathbf{k}_5 + \tilde{d}_6\mathbf{k}_6. \quad (\text{B.13})$$

In practice we form the single error estimate

$$\Delta_m = \max_j \left\{ \frac{|c_j - \tilde{c}_j|^2}{|c_j|^2} \right\}, \quad (\text{B.14})$$

Appendix B. Interaction-picture integration algorithms

where we only consider those mode elements c_j such that $|c_j|^2 > 10^{-3} \max\{|c_i|^2\}$, and thus either accept the step and increase Δt , or reject the step and decrease Δt , according to (B.11), with safety factor $S = 0.92$ [63].

Appendix C

Quadrature methods for the PGPE

In this appendix we show how Gaussian quadrature rules can be used to efficiently and accurately evaluate the integrals which appear in the spectral-Galerkin representation of the PGPE. We closely follow the development presented by Blakie [27] and Bradley [37, 34].

C.1 Gaussian quadrature rules

A Gaussian quadrature rule is a means of approximation an integral by a discrete sum, i.e.

$$\int_a^b dx W(x) f(x) \approx \sum_{j=1}^N w_j f(x_j), \quad (\text{C.1})$$

where $W(x)$ is some *weight function*, and the w_j are the quadrature *weights* corresponding to the *abscissas* x_j that define the appropriate *quadrature grid*, which is in general not evenly spaced. A particularly nice feature of such Gaussian quadrature rules is that in the case that $f(x)$ is a polynomial, the quadrature-rule order N can be chosen such that the approximate correspondence of equation (C.1) becomes *exact*. Such quadrature rules have deep connection with the theory of orthogonal polynomials: given a domain of integration $[a, b]$, weight function $W(x)$, and set of polynomials $\{p_j(\mathbf{x})\}$ which are orthonormal with respect to the weighted integral on this domain, i.e.

$$\langle p_i | p_j \rangle \equiv \int_a^b dx W(x) p_i(x) p_j(x) = \delta_{ij}, \quad (\text{C.2})$$

Appendix C. Quadrature methods for the PGPE

the N roots of $p_N(x)$ are precisely the N abscissas x_j for the quadrature rule on $[a, b]$ with weight $W(x)$. The corresponding weights are given by

$$w_j = \frac{\langle p_{N-1} | p_{N-1} \rangle}{p_{N-1}(x_j) p'_N(x_j)}, \quad (\text{C.3})$$

where $p'_N(x_j)$ is the derivative of $p_N(x)$ evaluated at x_j . For a quadrature rule of order N , the substitution (C.1) is then exact for all *general* polynomials of degree $2N - 1$ or less.

For example, for the particular weight function $W(x) = e^{-x^2}$ and domain of integration $(-\infty, \infty)$, the corresponding orthogonal polynomials are the *Hermite polynomials* defined by

$$H_0(x) = 1 \quad (\text{C.4a})$$

$$H_1(x) = 2x \quad (\text{C.4b})$$

$$H_{j+1}(x) = 2xH_j(x) - 2jH_{j-1}(x), \quad (\text{C.4c})$$

and the corresponding rule (C.1) is the Gauss-Hermite quadrature rule. The quadrature abscissas and corresponding weights for this rule are tabulated in (e.g.) [3].

C.2 Hermite quadratures for harmonic confinement

In this section we describe the Gauss-Hermite quadrature method for evaluating the (projected) nonlinear term in the nonrotating PGPE. This discussion closely follows that of Blakie [27].

C.2.1 Application to the PGPE

We begin with the dimensionless PGPE

$$i\partial_t\psi = H_0\psi + \mathcal{P} \{U_0|\psi|^2\psi\}, \quad (\text{C.5})$$

where the dimensionless base Hamiltonian

$$H_0 = \frac{-\nabla^2}{2} + \frac{1}{2} (\lambda_x^2 x^2 + \lambda_y^2 y^2 + \lambda_z^2 z^2). \quad (\text{C.6})$$

For simplicity, we follow Blakie [27] in considering the isotropic case $\lambda_x = \lambda_y = \lambda_z = 1$. The extension to anisotropic harmonic trapping is straightforward. As the Hamiltonian H_0 can be separated into additive parts corresponding to the single-particle degrees of

freedom in each of the spatial dimensions, i.e. $H_0 = H_{0x} + H_{0y} + H_{0z}$, the eigenstates of H_0 are separable, i.e. $\varphi_n(\mathbf{x}) = \phi_\alpha(x)\phi_\beta(y)\phi_\gamma(z)$, where (e.g.) the eigenstates $\phi_\alpha(x)$ of H_{0x} are such that

$$\left[-\frac{1}{2} \frac{d^2}{dx^2} + \frac{1}{2} x^2 \right] \phi_\alpha(x) = \epsilon_\alpha \phi_\alpha(x), \quad (\text{C.7})$$

where the eigenvalues $\epsilon_\alpha = \alpha + \frac{1}{2}$, so that the eigenvalues of H_0 are

$$\epsilon_n = \epsilon_\alpha + \epsilon_\beta + \epsilon_\gamma. \quad (\text{C.8})$$

The eigenfunctions themselves are of form

$$\phi_\alpha(x) = h_\alpha H_\alpha(x) e^{-x^2/2}, \quad (\text{C.9})$$

where the normalisation coefficient

$$h_\alpha = \frac{1}{\sqrt{2^\alpha \alpha! \sqrt{\pi}}}, \quad (\text{C.10})$$

and the $H_\alpha(x)$ are the Hermite polynomials (equations (C.4-C.4c)). The coherent region can thus be defined

$$\mathbf{C} \equiv \{\alpha, \beta, \gamma : \epsilon_\alpha + \epsilon_\beta + \epsilon_\gamma \leq E_R\}, \quad (\text{C.11})$$

and the maximum allowed value of each of α, β, γ is therefore

$$\alpha_{\max} = E_R - \frac{3}{2}, \quad (\text{C.12})$$

yielding $M = \alpha_{\max} + 1$ distinct eigenstates in each dimension. The equation (C.11) defines the coherent region \mathbf{C} as a pyramidal volume in $\alpha\beta\gamma$ space, which thus encloses a volume (i.e., number of modes) $\mathcal{M} \approx (1/6)M^3$. Our task in evaluating the nonlinear potential is therefore that of evaluating the integral

$$G_{\alpha\beta\gamma}[\psi(\mathbf{x})] = \int d\mathbf{x} \phi_\alpha^*(x) \phi_\beta^*(y) \phi_\gamma^*(z) |\psi(\mathbf{x})|^2 \psi(\mathbf{x}). \quad (\text{C.13})$$

The important point is that, given the restriction (C.11) on the single-particle modes from which $\psi(\mathbf{x})$ is constructed, we have

$$\psi(\mathbf{x}) = Q(x, y, z) e^{-(x^2+y^2+z^2)/2}, \quad (\text{C.14})$$

Appendix C. Quadrature methods for the PGPE

where $Q(x, y, z)$ is a polynomial

$$Q(x, y, z) = \sum_{\alpha\beta\gamma \in \mathbf{C}} c_{\alpha\beta\gamma} h_{\alpha} h_{\beta} h_{\gamma} H_{\alpha}(x) H_{\beta}(y) H_{\gamma}(z), \quad (\text{C.15})$$

of maximum degree $M - 1$ in each of x , y , and z . As a result, the nonlinear integral is of form

$$G_{\alpha\beta\gamma} = \int d\mathbf{x} e^{-2(x^2+y^2+z^2)} P_{\alpha\beta\gamma}(x, y, z), \quad (\text{C.16})$$

where

$$P_{\alpha\beta\gamma}(x, y, z) = h_{\alpha} h_{\beta} h_{\gamma} H_{\alpha}(x) H_{\beta}(y) H_{\gamma}(z) |Q(x, y, z)|^2 Q(x, y, z), \quad (\text{C.17})$$

is therefore of maximum degree $4(M - 1)$ in each of x , y , and z . Correspondingly, the integral $G_{\alpha\beta\gamma}$ is of the appropriate form for Gauss-Hermite quadrature, and can therefore be evaluated *exactly* using a three-dimensional spatial grid of size $2(M - 1) \times 2(M - 1) \times 2(M - 1)$ points, i.e.

$$G_{\alpha\beta\gamma} = \sum_{ijk} w_i w_j w_k P_{\alpha\beta\gamma}(x_i, x_j, x_k), \quad (\text{C.18})$$

where the x_i and w_i are respectively the *roots* and *weights* of the one-dimensional Gauss-Hermite quadrature, with weight function $w(x) = \exp(-2x^2)$ (see e.g., [3]).

The procedure for evaluating the projection of the nonlinear term is thus as follows:

1. Transform the field from the coefficient representation $\{c_{\alpha\beta\gamma}\}$ to a spatial representation on the *quadrature grid* $\{\mathbf{x}_{ijk}\}$, according to

$$\psi(\mathbf{x}_{ijk}, t) = \sum_{\alpha\beta\gamma \in \mathbf{C}} U_{i\alpha} U_{j\beta} U_{k\gamma} c_{\alpha\beta\gamma}(t), \quad (\text{C.19})$$

where the (precomputed) transformation matrices are simply the 1D basis states evaluated at the quadrature grid points, i.e.

$$U_{i\alpha} = \phi_{\alpha}(x_i). \quad (\text{C.20})$$

2. Construct the integrand on the quadrature grid, by dividing by the weight function and premultiplying by the weights

$$g(\mathbf{x}_{ijk}) = w_i w_j w_k e^{2|\mathbf{x}_{ijk}|^2} |\psi(\mathbf{x}_{ijk}, t)|^2 \psi(\mathbf{x}_{ijk}, t). \quad (\text{C.21})$$

3. Complete the integration by transforming back

$$G_{\alpha\beta\gamma} = \sum_{ijk} U_{i\alpha}^* U_{j\beta}^* U_{k\gamma}^* g(\mathbf{x}_{ijk}). \quad (\text{C.22})$$

C.3 Laguerre quadratures for the rotating system

In this section we describe the evaluation of the projected terms in the two-dimensional rotating-frame PGPE, including the evaluation of an elliptic perturbing-potential term. This discussion closely follows that of Bradley [34, 35]

C.3.1 Nonlinear term

We begin with the dimensionless PGPE

$$i\partial_t\psi = H_0\psi + \mathcal{P}\{U_{2D}|\psi|^2\psi\}, \quad (\text{C.23})$$

where the dimensionless rotating-frame base Hamiltonian

$$H_0 = -\frac{1}{2r}\frac{\partial}{\partial r}r\frac{\partial}{\partial r} - \frac{1}{2r^2}\frac{\partial^2}{\partial\theta^2} + \frac{r^2}{2} + i\Omega\frac{\partial}{\partial\theta}, \quad (\text{C.24})$$

with rotation frequency $0 \leq \Omega < 1$. The eigenstates of H_0 are

$$Y_{nl}(r, \theta) = \sqrt{\frac{n!}{\pi(n+|l|)!}} e^{il\theta} r^{|l|} e^{-r^2/2} L_n^{|l|}(r^2), \quad (\text{C.25})$$

where $L_n^m(x)$ denotes the *associated Laguerre polynomials*, defined by

$$L_0^{(\alpha)}(x) = 1 \quad (\text{C.26a})$$

$$L_1^{(\alpha)}(x) = -x + \alpha + 1 \quad (\text{C.26b})$$

$$nL_n^{(\alpha)}(x) = (2n + \alpha - 1 - x)L_{n-1}^{(\alpha)}(x) - (n + \alpha - 1)L_{n-2}^{(\alpha)}(x). \quad (\text{C.26c})$$

The eigenvalues of the modes $Y_{nl}(r, \theta)$ are

$$\omega_{nl} = 2n + |l| - \Omega l + 1. \quad (\text{C.27})$$

Introducing the quantity \bar{N} such that $E_R = \hbar\omega_r(\bar{N} + 1)$, the coherent region can be defined by

$$\mathbf{C} = \{n, l : 2n + |l| - \Omega l \leq \bar{N}\}. \quad (\text{C.28})$$

Appendix C. Quadrature methods for the PGPE

The ranges of allowed quantum numbers (n, l) are thus

$$0 \leq n \leq \left\lfloor \frac{\bar{N}}{2} \right\rfloor, \quad (\text{C.29})$$

$$-l_-(n) \equiv \left\lfloor \frac{\bar{N} - 2n}{1 + \Omega} \right\rfloor \leq l \leq \left\lfloor \frac{\bar{N} - 2n}{1 - \Omega} \right\rfloor \equiv l_+(n), \quad (\text{C.30})$$

where $\lfloor \dots \rfloor$ represents the *floor function*. The projected wave function is therefore

$$\psi(r, \theta) = \sum_{n=0}^{\lfloor \bar{N}/2 \rfloor} \sum_{l=-l_-(n)}^{l_+(n)} c_{nl} Y_{nl}(r, \theta). \quad (\text{C.31})$$

Our task in evaluating the nonlinear potential is thus that of evaluating

$$F_{nl}[\psi(r, \theta)] = \int_0^{2\pi} d\theta \int_0^\infty dr r Y_{nl}^*(r, \theta) |\psi(r, \theta)|^2 \psi(r, \theta), \quad (\text{C.32})$$

A change of variables $x = r^2$ gives

$$F_{nl} = \frac{1}{2\pi} \int_0^{2\pi} d\theta e^{-il\theta} \int_0^\infty dx |\psi(\sqrt{x}, \theta)|^2 \psi(\sqrt{x}, \theta) \Phi_{nl}(x) \pi, \quad (\text{C.33})$$

where

$$\Phi_{nl}(x) = \sqrt{\frac{n!}{\pi(n + |l|)!}} e^{-x/2} x^{|l|/2} L_n^{|l|}(x), \quad (\text{C.34})$$

and

$$\psi(\sqrt{x}, \theta) = \sum_{nl} c_{nl} e^{il\theta} \Phi_{nl}(x). \quad (\text{C.35})$$

The integral to compute is thus of form

$$I = \frac{1}{2\pi} \int_0^{2\pi} d\theta e^{-il\theta} \int_0^\infty dx e^{-2x} Q(x, \theta), \quad (\text{C.36})$$

where $Q(x, \theta)$ is a polynomial in x and $e^{i\theta}$ of a maximum order constrained by the cut-off. As ψ is a (Laurent) polynomial in $e^{i\theta}$ of maximum positive order $\bar{l}_+ = l_+(0)$ and maximum *negative* order $\bar{l}_- = l_-(0)$, the function $Q(\theta)$ will in general contain all powers $(e^{i\theta})^q$ where $-2(\bar{l}_- + \bar{l}_+) \leq q \leq 2(\bar{l}_+ + \bar{l}_-)$. The angular integral can thus be evaluated *exactly* as a discrete Fourier transform (DFT), discretising θ as $\theta_j = j\Delta = j2\pi/N_\theta$, for

$j = 0, 1, \dots, N_\theta - 1$, where $N_\theta \equiv 2(\bar{l}_+ + \bar{l}_-) + 1$, to give

$$\int_0^{2\pi} d\theta e^{-il\theta} Q(x, \theta) = \frac{2\pi}{N_\theta} \sum_{j=0}^{N_\theta-1} e^{-i2\pi jl/N_\theta} Q(x, \theta_j). \quad (\text{C.37})$$

To perform the radial integral, we note that the radial part of ψ is of the form $e^{-x/2} P(x)$, where $P(x)$ is a polynomial of at most order $\bar{l}_+/2$, and the radial integrand is thus of maximum order $2\bar{l}_+ = 2[\bar{N}/(1 - \Omega)]$. The radial integral is therefore of form

$$I_r = \int_0^\infty dx e^{-2x} Q(x, \theta), \quad (\text{C.38})$$

where $Q(x)$ is of order at most $2[\bar{N}/(1 - \Omega)]$. The integral I_r is thus of the appropriate form for *Gauss-Laguerre* quadrature, and can therefore be evaluated *exactly* on a grid with $N_x \equiv [\bar{N}/(1 - \Omega)] + 1$ points x_k , which are the roots of the (simple) Laguerre polynomial $L_{N_x}(x)$ with corresponding weights w_k , which are tabulated in (e.g.) reference [56], and we have

$$I_r = \frac{1}{2} \sum_{k=1}^{N_x} w_k Q(x_k/2, \theta). \quad (\text{C.39})$$

Combining the two quadrature rules we find

$$\begin{aligned} F_{nl} &= \sum_{k=1}^{N_x} \tilde{w}_k \Phi_{nl}(x_k/2) \frac{1}{N_\theta} \sum_{j=0}^{N_\theta-1} e^{-i2\pi jl/N_\theta} \\ &\times \left| \psi(\sqrt{x_k/2}, \theta_j) \right|^2 \psi(\sqrt{x_k/2}, \theta_j), \end{aligned} \quad (\text{C.40})$$

where we have grouped the weights as $\tilde{w}_k = \pi w_k e^{x_k/2}$. Once the radial parts of the basis modes are precomputed at the quadrature abscissas

$$P_{kn}^{[l]} = \Phi_{n|l|}(x_k/2), \quad (\text{C.41})$$

we have the following procedure:

1. Perform the radial part of the transformation from the coefficient representation $\{c_{nl}\}$ to the spatial representation

$$\chi_{kl} = \sum_{n=0}^{[\bar{N}/2]} P_{kn}^{[l]} c_{nl}. \quad (\text{C.42})$$

Appendix C. Quadrature methods for the PGPE

2. Construct the position-space field at the quadrature points Ψ_{kj} using the DFT, after zero padding χ_{kl} to length N_θ in the l index, $\tilde{\chi}_{kl} \equiv \text{pad}_l(\chi_{kl}, N_\theta)$,

$$\Psi_{kj} = \psi(\sqrt{x_k/2}, \theta_j) = \sum_{l=0}^{N_\theta-1} \tilde{\chi}_{kl} e^{i2\pi jl/N_\theta}. \quad (\text{C.43})$$

3. Form the integrand

$$\Xi_{kj} = |\Psi_{kj}|^2 \Psi_{kj}. \quad (\text{C.44})$$

4. Perform the azimuthal part of the inverse transform by DFT

$$\Theta_{kl} = \frac{1}{N_\theta} \sum_{j=0}^{N_\theta-1} \Xi_{kj} e^{-i2\pi jl/N_\theta}. \quad (\text{C.45})$$

5. Perform the radial part of the inverse transform by Gauss-Laguerre quadrature

$$F_{nl} = \sum_{k=1}^{N_x} \tilde{w}_k \Theta_{kl} P_{kn}^{|l|}. \quad (\text{C.46})$$

C.3.2 Elliptic trap anisotropy

The general form of the elliptic potential anisotropy we need to compute (chapters 7 and 9) is

$$V_\epsilon(r, \theta) = -\frac{\epsilon}{2} r^2 \cos(2\theta). \quad (\text{C.47})$$

The corresponding term in the PGPE is

$$H_{nl} = -\frac{\epsilon}{2} \int_0^{2\pi} d\theta \int_0^\infty dr r^3 Y_{nl}^*(r, \theta) |\psi(r, \theta)|^2 \psi(r, \theta); \quad (\text{C.48})$$

making the change of variables $x = r^2$, we have

$$\begin{aligned} H_{nl} &= -\frac{\epsilon}{4\pi} \int_0^{2\pi} d\theta e^{-il\theta} \cos(2\theta) \\ &\quad \times \int_0^\infty dx x \Phi_{nl}(x) \psi(\sqrt{x}, \theta) \pi. \end{aligned} \quad (\text{C.49})$$

The integral to evaluate is therefore of the form

$$H_{nl} = \frac{1}{4\pi} \int_0^{2\pi} d\theta e^{-il\theta} \cos(2\theta) \int_0^\infty dx e^{-x} Q(x, \theta), \quad (\text{C.50})$$

where $Q(x, \theta)$ is a polynomial in x and $e^{i\theta}$ of order determined by the cutoff. Due to the constraint on the powers of $e^{i\theta}$ that appear in $\psi(\sqrt{x}, \theta)$, the function $Q(x, \theta)$ will in general contain all powers $(e^{i\theta})^q$ where $-\bar{l}_- \leq q \leq \bar{l}_+$. As a result the total radial integrand will in general contain all powers $-(\bar{l}_- + \bar{l}_+ + 2) \leq q \leq (\bar{l}_- + \bar{l}_+ + 2)$, and the integral can thus be carried out on an azimuthal grid of $N_\theta^{(2)} = \bar{l}_- + \bar{l}_+ + 3$ points, discretising θ as $\theta_j^{(2)} = j\Delta\theta^{(2)} = j2\pi/N_\theta^{(2)}$, $0 \leq j \leq N_\theta^{(2)} - 1$. The azimuthal integral can thus be evaluated as the DFT

$$\int_0^{2\pi} d\theta e^{-il\theta} \cos(2\theta) Q(x, \theta) = \frac{2\pi}{N_\theta^{(2)}} \sum_{j=0}^{N_\theta^{(2)}-1} e^{-i2\pi jl/N_\theta^{(2)}} \cos(2\theta_j^{(2)}) Q(x, \theta_j^{(2)}). \quad (\text{C.51})$$

To perform the radial integral, we note that the radial part is of the form $e^{-x}Q(x, \theta)$, where Q is of maximum order \bar{l}_+ in x . The radial integral can therefore be computed exactly using a Gauss-Laguerre quadrature rule of order $N_x^{(2)} \equiv [\bar{l}_+/2] + 1$. The radial integral is thus

$$\int_0^\infty dx e^{-x} Q(x, \theta) = \sum_{k=1}^{N_x^{(2)}} w_k^{(2)} Q(x_k^{(2)}, \theta), \quad (\text{C.52})$$

where the $x_k^{(2)}$ are the roots of $L_{N_x^{(2)}}(x)$, and $w_k^{(2)}$ are the corresponding weights [56]. Combining the two quadrature rules we thus find

$$\begin{aligned} H_{nl} = & -\frac{\epsilon}{2} \sum_{k=1}^{N_x^{(2)}} \tilde{w}_k^{(2)} \Phi_{nl}(x_k^{(2)}) \frac{1}{N_\theta^{(2)}} \sum_{j=0}^{N_\theta^{(2)}-1} e^{-i2\pi jl/N_\theta^{(2)}} \\ & \times \cos(4\pi jl/N_\theta^{(2)}) \psi(\sqrt{x_k^{(2)}}, \theta_j^{(2)}), \end{aligned} \quad (\text{C.53})$$

where we have grouped the weights as $\tilde{w}_k^{(2)} = \pi w_k^{(2)} e^{x_k^{(2)}/2}$. Once the modes are precomputed at the quadrature abscissas

$$P_{kn}^{[l]} = \Phi_{n|l]}(x_k^{(2)}), \quad (\text{C.54})$$

we have the following procedure:

Appendix C. Quadrature methods for the PGPE

1. Perform the radial part of the transformation

$$\chi_{kl} = \sum_{n=0}^{[\bar{N}/2]} P_{kn}^{[l]} c_{nl}. \quad (\text{C.55})$$

2. Construct the position-space field at the quadrature points, after zero padding χ_{kl} to length $N_\theta^{(2)}$ in the l index, $\tilde{\chi}_{kl} \equiv \text{pad}_l(\chi_{kl}, N_\theta^{(2)})$,

$$\Psi_{kj} = \sum_{l=0}^{N_\theta^{(2)}-1} e^{i2\pi jl/N_\theta^{(2)}} \tilde{\chi}_{kl}. \quad (\text{C.56})$$

3. Apply the perturbing potential to form the integrand

$$\Gamma_{kj} = \Psi_{kj} x_k^{(2)} \cos(2\theta_j^{(2)}). \quad (\text{C.57})$$

4. Perform the azimuthal part of the inverse transform by DFT

$$\Theta_{kl} = \frac{2}{N_\theta^{(2)}} \sum_{j=0}^{N_\theta^{(2)}-1} \Gamma_{kj} e^{-i2\pi jl/N_\theta^{(2)}}. \quad (\text{C.58})$$

5. Perform the radial part of the inverse transform to evaluate the Gauss-Laguerre quadrature

$$H_{nl} = -\frac{\epsilon}{2} \sum_{k=1}^{N_x^{(2)}} \tilde{w}_k^{(2)} \Theta_{kl} P_{kn}^{[l]}. \quad (\text{C.59})$$

The orders of quadrature rules used here are significantly smaller than those used for the evaluation of the nonlinear term, and consequently the computational load increase due to the inclusion of the trap anisotropy over that of the base method is of order 25 – 50%, depending on the rotation rate and cutoff height (see section 4.2).

Appendix D

Semiclassical fitting function

In order to estimate the thermodynamic parameters of the classical field, we assume an approximate semiclassical description for the thermal atoms. The description consists of two elements: a classical-field distribution in the semiclassical phase-space coordinates (\mathbf{x} and \mathbf{p}), and the specification of an appropriate energy cutoff. In order to construct this distribution, we must take into account two generally disparate rotations: the (quasi-)equilibrium rotation of the thermal atoms (angular velocity Ω_{th}) and the rotation of the projector defining the cutoff (angular velocity Ω_{p}). The classical-field distribution takes the general form (cf. equation (3.31))

$$F_c(\mathbf{x}, \mathbf{p}) = \frac{1}{h^2} \frac{k_B T}{\epsilon(\mathbf{x}, \mathbf{p}) - \mu}, \quad (\text{D.1})$$

where the semiclassical Hartree-Fock energy of the thermal atoms (rotating at Ω_{th}) is

$$\epsilon(\mathbf{x}, \mathbf{p}) = \frac{p^2}{2m} + \frac{m\omega_r^2 r^2}{2} - \Omega_{\text{th}}(xp_y - yp_x) + 2U_{2\text{D}}n(r), \quad (\text{D.2})$$

(cf. [251]) and where $n(r)$ is the total density of the atomic field, which we will take to be circularly symmetric. The application of the projector \mathcal{P} defined in a rotating frame enforces a cutoff in the energy, which delimits the accessible region of phase space occupied according to equation (D.1). This region is defined by

$$\frac{p^2}{2m} + \frac{m\omega_r^2 r^2}{2} - \Omega_{\text{p}}(xp_y - yp_x) + 2U_{2\text{D}}n(r) \leq E_R. \quad (\text{D.3})$$

The density of thermal atoms at any spatial position \mathbf{x} is obtained by integrating equation (D.1) over the momentum, i.e. $n(\mathbf{x}) = \int_{\mathcal{D}} d\mathbf{p} F_c(\mathbf{x}, \mathbf{p})$, where the domain \mathcal{D} in momentum space is determined by equation (D.3), and is itself a function of the position-space coordinate \mathbf{x} . To perform this integration we make the transformation to the

Appendix D. Semiclassical fitting function

kinematic momentum measured in the frame of the cloud $\mathbf{P} = \mathbf{p} - m\boldsymbol{\Omega}_{\text{th}} \times \mathbf{x}$ [251, 156]. In this representation the integrand becomes circularly symmetric in \mathbf{P} , i.e., we now have $F'_c(\mathbf{x}, \mathbf{P}) = (1/h^2)k_B T/[\epsilon(\mathbf{x}, \mathbf{P}) - \mu]$, where

$$\epsilon(\mathbf{x}, \mathbf{P}) = \frac{P^2}{2m} + \frac{m}{2} (\omega_r^2 - \Omega_{\text{th}}^2) r^2 + 2U_{2D}n(r). \quad (\text{D.4})$$

In these momentum coordinates the domain \mathcal{D} is circular, however its centre is in general displaced from the origin $\mathbf{P} = \mathbf{0}$. After some simple algebra we find that the domain of integration is

$$\begin{aligned} \mathcal{D}(\mathbf{x}) : \quad & \frac{1}{2m} \left\{ [P_x + m(\Omega_p - \Omega_{\text{th}})y]^2 + [P_y - m(\Omega_p - \Omega_{\text{th}})x]^2 \right\} \\ & \leq E_R - \frac{m}{2} (\omega_r^2 - \Omega_p^2) r^2 - 2U_{2D}n(r). \end{aligned} \quad (\text{D.5})$$

As the integrand is circularly symmetric in \mathbf{P} , the result of the integral depends only on the radius r of the chosen point in position space, as azimuthal rotations of the coordinate \mathbf{x} simply induce rotations of $\mathcal{D}(\mathbf{x})$ about the origin of momentum space. We therefore consider without loss of generality $(x, y) = (0, -r)$. The domain \mathcal{D} thus becomes a circle centred on $P_x = m(\Omega_p - \Omega_{\text{th}})r \equiv a$, of radius $b \equiv \sqrt{2m[E_R - (m/2)(\omega_r^2 - \Omega_p^2)r^2 - 2U_{2D}n(r)]}$. Depending on the value of r , we may have $a \leq b$ or $b > a$. We adopt polar coordinates (P, ϕ) , and find for the angle $\Delta\phi$ subtended by \mathcal{D} at a given radius P

$$\Delta\phi(P) = \begin{cases} 0 & P < a - b \quad (\text{when } a > b) \\ 2\pi & P < b - a \quad (\text{when } a < b) \\ 2 \cos^{-1} \left(\frac{P^2 - b^2 + a^2}{2aP} \right) & |a - b| < P < a + b \\ 0 & P > a + b \end{cases}. \quad (\text{D.6})$$

Integrating in polar coordinates we obtain finally

$$\begin{aligned} n(r) &= \frac{2\pi m k_B T}{h^2} \int_0^\infty \frac{P dP \Delta\phi(P)}{\frac{P^2}{2m} + \frac{m\omega_r^2 r^2}{2} + 2U_{2D}n(r) - \mu} \\ &= \frac{1}{\lambda_{\text{dB}}^2} \left[I_1(r; \Omega_{\text{th}}, \Omega_p) + I_2(r; \Omega_{\text{th}}, \Omega_p) \right], \end{aligned} \quad (\text{D.7})$$

where

$$I_1(r) = \Theta(b - a) \ln \left[\frac{\frac{(b-a)^2}{2m} + \frac{m}{2} (\omega_r^2 - \Omega_{\text{th}}^2) r^2 + 2U_{2D}n(r) - \mu}{\frac{m}{2} (\omega_r^2 - \Omega_{\text{th}}^2) r^2 + 2U_{2D}n(r) - \mu} \right], \quad (\text{D.8})$$

and

$$I_2(r) = \frac{1}{\pi} \int_{|b-a|}^{a+b} \frac{P dP \cos^{-1} \left(\frac{\frac{p^2}{2m} + \frac{m}{2} [\omega_r^2 - \Omega_p^2 + (\Omega_p - \Omega_{\text{th}})^2] r^2 - E_R}{(\Omega_p - \Omega_{\text{th}}) r P} \right)}{\frac{p^2}{2m} + \frac{m}{2} (\omega_r^2 - \Omega_{\text{th}}^2) r^2 - \mu}, \quad (\text{D.9})$$

and the de Broglie wavelength $\lambda_{\text{dB}} = \sqrt{2\pi\hbar^2/mk_{\text{B}}T}$. The integral I_2 must be evaluated numerically. In the limit $\Omega_{\text{th}} \rightarrow \Omega_p$ the integral I_2 vanishes, and we obtain the form

$$n(r) = \frac{1}{\lambda_{\text{dB}}^2} \ln \left[\frac{E_R - \mu}{(\omega_r^2 - \Omega_{\text{th}}^2) r^2 + 2U_{2\text{D}} n(r) - \mu} \right], \quad (\text{D.10})$$

employed in chapter 9.

References

- [1] J. R. Abo-Shaeer, C. Raman, and W. Ketterle, “Formation and Decay of Vortex Lattices in Bose-Einstein Condensates at Finite Temperatures,” *Phys. Rev. Lett.* **88**, 070409 (2002).
- [2] J. R. Abo-Shaeer, C. Raman, J. M. Vogels, and W. Ketterle, “Observation of Vortex Lattices in Bose-Einstein Condensates,” *Science* **292**, 476 (2001).
- [3] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions* (Dover, New York, 1972).
- [4] A. A. Abrikosov, “On the Magnetic Properties of Superconductors of the Second Group,” *Sov. Phys. JETP* **5**, 1174 (1957).
- [5] A. A. Abrikosov, L. P. Gor’kov, and I. E. Dzyaloshinskii, *Methods of Quantum Field Theory in Statistical Physics* (Prentice-Hall, Englewood Cliffs, New Jersey, 1963).
- [6] M. F. Andersen, C. Ryu, P. Cladé, V. Natarajan, A. Vaziri, K. Helmerson, and W. D. Phillips, “Quantized Rotation of Atoms from Photons with Orbital Angular Momentum,” *Phys. Rev. Lett.* **97**, 170406 (2006).
- [7] B. P. Anderson, P. C. Haljan, C. E. Wieman, and E. A. Cornell, “Vortex Precession in Bose-Einstein Condensates: Observations with Filled and Empty Cores,” *Phys. Rev. Lett.* **85**, 2857 (2000).
- [8] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor,” *Science* **269**, 198 (1995).
- [9] P. W. Anderson, “Two new vortex liquids,” *Nature Phys.* **3**, 160 (2007).
- [10] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Observation of Interference Between Two Bose Condensates,” *Science* **275**, 637 (1997).

References

- [11] J. R. Anglin, “Local Vortex Generation and the Surface Mode Spectrum of Large Bose-Einstein Condensates,” *Phys. Rev. Lett.* **87**, 240401 (2001).
- [12] J. R. Anglin, “Second-quantized Landau-Zener theory for dynamical instabilities,” *Phys. Rev. A* **67**, 051601(R) (2003).
- [13] P. Ao and D. J. Thouless, “Berry’s phase and the Magnus force for a vortex line in a superconductor,” *Phys. Rev. Lett.* **70**, 2158 (1993).
- [14] H. Aref, “Integrable, Chaotic, and Turbulent Vortex Motion in Two-Dimensional Flows,” *Annu. Rev. Fluid Mech.* **15**, 345 (1983).
- [15] D. P. Arovas and J. A. Freire, “Dynamical vortices in superfluid films,” *Phys. Rev. B* **55**, 1068 (1997).
- [16] V. Bagnato, D. E. Pritchard, and D. Kleppner, “Bose-Einstein condensation in an external potential,” *Phys. Rev. A* **35**, 4354 (1987).
- [17] R. J. Ballagh, “Computational methods for nonlinear partial differential equations,” <http://www.physics.otago.ac.nz/nx/jdc/computation-lectures.html> (2000).
- [18] H. P. Baltes and E. R. Hilf, *Spectra of Finite Systems* (Bibliographisches Institut, Zurich, 1976).
- [19] C. F. Barenghi, R. J. Donnelly, and W. F. Vinen, *Quantized Vortex Dynamics and Superfluid Turbulence* (Springer-Verlag, Berlin, 2001).
- [20] G. Baym, C. Pethick, and D. Pines, “Superfluidity in Neutron Stars,” *Nature* **224**, 673 (1969).
- [21] T. Bergeman, D. L. Feder, N. L. Balazs, and B. I. Schneider, “Bose condensates in a harmonic trap near the critical temperature,” *Phys. Rev. A* **61**, 063605 (2000).
- [22] N. G. Berloff and A. J. Youd, “Dissipative Dynamics of Superfluid Vortices at Nonzero Temperatures,” *Phys. Rev. Lett.* **99**, 145301 (2007).
- [23] A. Bezett and P. B. Blakie, “Critical properties of a trapped interacting Bose gas,” *Phys. Rev. A* **79**, 033611 (2009).
- [24] A. Bezett, E. Toth, and P. B. Blakie, “Two-point correlations of a trapped interacting Bose gas at finite temperature,” *Phys. Rev. A* **77**, 023602 (2008).

-
- [25] J.-P. Blaizot and G. Ripka, *Quantum Theory of Finite Systems* (MIT Press, Cambridge, Massachusetts, 1986).
- [26] P. B. Blakie, *Optical Manipulation of Bose-Einstein Condensates*, Ph.D. thesis, University of Otago (2001).
- [27] P. B. Blakie, “Numerical method for evolving the projected Gross-Pitaevskii equation,” *Phys. Rev. E* **78**, 026704 (2008).
- [28] P. B. Blakie, A. S. Bradley, M. J. Davis, R. J. Ballagh, and C. W. Gardiner, “Dynamics and statistical mechanics of ultra-cold Bose gases using c-field techniques,” *Adv. Phys.* **57**, 363 (2008).
- [29] P. B. Blakie and M. J. Davis, “Projected Gross-Pitaevskii equation for harmonically confined Bose gases at finite temperature,” *Phys. Rev. A* **72**, 063608 (2005).
- [30] P. B. Blakie and M. J. Davis, “Classical region of a trapped Bose gas,” *J. Phys. B* **40**, 2043 (2007).
- [31] G. Blatter, M. V. Feigel'man, V. B. Geshkenbein, A. I. Larkin, and V. M. Vinokur, “Vortices in high-temperature superconductors,” *Rev. Mod. Phys.* **66**, 1125 (1994).
- [32] I. Bloch, J. Dalibard, and W. Zwerger, “Many-body physics with ultracold gases,” *Rev. Mod. Phys.* **80**, 885 (2008).
- [33] N. N. Bogoliubov, “On the theory of superfluidity,” *J. Phys. (U.S.S.R.)* **11**, 23 (1947).
- [34] A. S. Bradley, *Theoretical investigations of trapped ultra-cold Bose gases: Rotating Bose-Einstein condensation and Anomalous Segregation*, Ph.D. thesis, Victoria University of Wellington (2005).
- [35] A. S. Bradley, “Numerical scheme for rotating elliptic potential,” unpublished personal papers (2006).
- [36] A. S. Bradley, P. B. Blakie, and C. W. Gardiner, “Properties of the stochastic Gross-Pitaevskii equation: finite temperature Ehrenfest relations and the optimal plane wave representation,” *J. Phys. B* **38**, 4259 (2005).
- [37] A. S. Bradley, C. W. Gardiner, and M. J. Davis, “Bose-Einstein condensation from a rotating thermal cloud: Vortex nucleation and lattice formation,” *Phys. Rev. A* **77**, 033616 (2008).

References

- [38] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, “Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions,” *Phys. Rev. Lett.* **75**, 1687 (1995).
- [39] B. H. Bransden and C. J. Joachain, *Quantum Mechanics*, 2nd edition (Pearson Education, Essex, 2000).
- [40] M. Brewczyk, P. Borowski, M. Gajda, and K. Rzążewski, “Temperature-dependent Bogoliubov approximation in the classical field approach to weakly interacting Bose gases,” *J. Phys. B* **37**, 2725 (2004).
- [41] S. Burger, K. Bongs, S. Dettmer, W. Ertmer, K. Sengstock, A. Sanpera, G. V. Shlyapnikov, and M. Lewenstein, “Dark Solitons in Bose-Einstein Condensates,” *Phys. Rev. Lett.* **83**, 5198 (1999).
- [42] D. A. Butts and D. S. Rokhsar, “Predicted signatures of rotating Bose-Einstein condensates,” *Nature* **397**, 327 (1999).
- [43] L. J. Campbell and R. M. Ziff, “Vortex patterns and energies in a rotating superfluid,” *Phys. Rev. B* **20**, 1886 (1979).
- [44] B. M. Caradoc-Davies, *Vortex Dynamics in Bose-Einstein Condensates*, Ph.D. thesis, University of Otago (2000).
- [45] B. M. Caradoc-Davies, R. J. Ballagh, and K. Burnett, “Coherent Dynamics of Vortex Formation in Trapped Bose-Einstein Condensates,” *Phys. Rev. Lett.* **83**, 895 (1999).
- [46] H. J. Carmichael, *Statistical Methods in Quantum Optics 1: Master Equations and Fokker-Planck Equations* (Springer-Verlag, Berlin, 1999).
- [47] Y. Castin, “Bose-Einstein condensates in atomic Gases: simple theoretical results,” in *Coherent Atomic Matter Waves*, edited by R. Kaiser, C. Westbrook, and F. David, Proceedings of the Les Houches Summer School of Theoretical Physics, Session LXXII (EDP Sciences and Springer-Verlag, Berlin, 2001).
- [48] Y. Castin and R. Dum, “Bose-Einstein Condensates in Time Dependent Traps,” *Phys. Rev. Lett.* **77**, 5315 (1996).
- [49] Y. Castin and R. Dum, “Instability and Depletion of an Excited Bose-Einstein Condensate in a Trap,” *Phys. Rev. Lett.* **79**, 3553 (1997).

-
- [50] Y. Castin and R. Dum, “Low-temperature Bose-Einstein condensates in time-dependent traps: Beyond the $U(1)$ symmetry-breaking approach,” *Phys. Rev. A* **57**, 3008 (1998).
- [51] Y. Castin and C. Herzog, “Bose-Einstein condensates in symmetry breaking states,” *C. R. Acad. Sci., Ser. IV: Phys., Astrophys.* **2**, 419 (2001).
- [52] F. Chevy, K. W. Madison, V. Bretin, and J. Dalibard, “Formation of quantized vortices in a gaseous Bose-Einstein condensate,” in *Trapped Particles and Fundamental Physics*, edited by S. N. Atutov, R. Calabrese, and L. Moi, Proceedings of the NATO Advanced Study Institute, Les Houches (Kluwer Academic, Dordrecht, 2002).
- [53] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, “Feshbach resonances in ultracold gases,” *Rev. Mod. Phys.* **82**, 1225 (2010).
- [54] S. Choi, S. A. Morgan, and K. Burnett, “Phenomenological damping in trapped atomic Bose-Einstein condensates,” *Phys. Rev. A* **57**, 4057 (1998).
- [55] P. Cladé, C. Ryu, A. Ramanathan, K. Helmerson, and W. D. Phillips, “Observation of a 2D Bose Gas: From Thermal to Quasicondensate to Superfluid,” *Phys. Rev. Lett.* **102**, 170401 (2009).
- [56] A. M. Cohen, J. F. Cutts, R. Fielder, D. E. Jones, J. Ribbens, and E. Stuart, *Numerical Analysis* (McGraw-Hill, London, 1973).
- [57] C. Connaughton, C. Josserand, A. Picozzi, Y. Pomeau, and S. Rica, “Condensation of Classical Nonlinear Waves,” *Phys. Rev. Lett.* **95**, 263901 (2005).
- [58] D. Dagnino, N. Barberán, M. Lewenstein, and J. Dalibard, “Vortex nucleation as a case study of symmetry breaking in quantum systems,” *Nature Phys.* **5**, 431 (2009).
- [59] F. Dalfovo, S. Giorgini, M. Guilleumas, L. Pitaevskii, and S. Stringari, “Collective and single-particle excitations of a trapped Bose gas,” *Phys. Rev. A* **56**, 3840 (1997).
- [60] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, “Theory of Bose-Einstein condensation in trapped gases,” *Rev. Mod. Phys.* **71**, 463 (1999).
- [61] F. Dalfovo and S. Stringari, “Bosons in anisotropic traps: Ground state and vortices,” *Phys. Rev. A* **53**, 2477 (1996).

References

- [62] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose-Einstein Condensation in a Gas of Sodium Atoms,” *Phys. Rev. Lett.* **75**, 3969 (1995).
- [63] M. J. Davis, *Dynamics of Bose-Einstein Condensation*, D.Phil. thesis, University of Oxford (2001).
- [64] M. J. Davis and P. B. Blakie, “Calculation of the microcanonical temperature for the classical Bose field,” *J. Phys. A* **38**, 10259 (2005).
- [65] M. J. Davis and P. B. Blakie, “Critical Temperature of a Trapped Bose Gas: Comparison of Theory and Experiment,” *Phys. Rev. Lett.* **96**, 060404 (2006).
- [66] M. J. Davis and S. A. Morgan, “Microcanonical temperature for a classical field: Application to Bose-Einstein condensation,” *Phys. Rev. A* **68**, 053615 (2003).
- [67] M. J. Davis, S. A. Morgan, and K. Burnett, “Simulations of Bose Fields at Finite Temperature,” *Phys. Rev. Lett.* **87**, 160402 (2001).
- [68] M. J. Davis, S. A. Morgan, and K. Burnett, “Simulations of thermal Bose fields in the classical limit,” *Phys. Rev. A* **66**, 053618 (2002).
- [69] B. DeMarco and D. S. Jin, “Onset of Fermi Degeneracy in a Trapped Atomic Gas,” *Science* **285**, 1703 (1999).
- [70] L. Deng, E. W. Hagley, J. Wen, M. Trippenbach, Y. Band, P. S. Julienne, J. E. Simsarian, K. Helmerson, S. L. Rolston, and W. D. Phillips, “Four-wave mixing with matter waves,” *Nature* **398**, 218 (1999).
- [71] C. M. Dion and E. Cancès, “Spectral method for the time-dependent Gross-Pitaevskii equation with a harmonic trap,” *Phys. Rev. E* **67**, 046706 (2003).
- [72] J. F. Dobson, “Harmonic-Potential Theorem: Implications for Approximate Many-Body Theories,” *Phys. Rev. Lett.* **73**, 2244 (1994).
- [73] R. J. Dodd, K. Burnett, M. Edwards, and C. W. Clark, “Excitation spectroscopy of vortex states in dilute Bose-Einstein condensed gases,” *Phys. Rev. A* **56**, 587 (1997).
- [74] R. J. Dodd, K. Burnett, M. Edwards, and C. W. Clark, “Trapped Bose-Einstein condensates at finite temperature: a two-gas model,” *Acta Phys. Pol. A* **93**, 45 (1998).

-
- [75] R. J. Donnelly, *Quantized Vortices in Helium II* (Cambridge University Press, Cambridge, 1991).
- [76] P. D. Drummond, T. Vaughan, J. Corney, G. Leuchs, and P. Deuar, “Coherence and Correlations in Atom-Lasers,” in *Conference on Coherence and Quantum Optics*, OSA Technical Digest (CD) (Optical Society of America, 2007), paper CWA1.
- [77] J. R. Ensher, D. S. Jin, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Bose-Einstein Condensation in a Dilute Gas: Measurement of Energy and Ground-State Occupation,” *Phys. Rev. Lett.* **77**, 4984 (1996).
- [78] J. Esteve, J.-B. Trebbia, T. Schumm, A. Aspect, C. I. Westbrook, and I. Bouchoule, “Observations of Density Fluctuations in an Elongated Bose Gas: Ideal Gas and Quasicondensate Regimes,” *Phys. Rev. Lett.* **96**, 130403 (2006).
- [79] D. L. Feder and C. W. Clark, “Superfluid-to-Solid Crossover in a Rotating Bose-Einstein Condensate,” *Phys. Rev. Lett.* **87**, 190401 (2001).
- [80] D. L. Feder, A. A. Svidzinsky, A. L. Fetter, and C. W. Clark, “Anomalous Modes Drive Vortex Dynamics in Confined Bose-Einstein Condensates,” *Phys. Rev. Lett.* **86**, 564 (2001).
- [81] P. O. Fedichev and G. V. Shlyapnikov, “Dissipative dynamics of a vortex state in a trapped Bose-condensed gas,” *Phys. Rev. A* **60**, R1779 (1999).
- [82] P. O. Fedichev, G. V. Shlyapnikov, and J. T. M. Walraven, “Damping of Low-Energy Excitations of a Trapped Bose-Einstein Condensate at Finite Temperatures,” *Phys. Rev. Lett.* **80**, 2269 (1998).
- [83] M. D. Feit, J. A. Fleck, Jr., and A. Steiger, “Solution of the Schrödinger Equation by a Spectral Method,” *J. Comput. Phys.* **47**, 412 (1982).
- [84] E. Fermi, J. Pasta, and S. Ulam, in *Collected Papers of Enrico Fermi*, edited by E. Segré, volume 2 (University of Chicago Press, Chicago, 1965).
- [85] A. L. Fetter, “Core Structure of a Quantized Vortex,” *Phys. Rev. Lett.* **27**, 986 (1971).
- [86] A. L. Fetter, “Theory of a dilute low-temperature trapped Bose condensate,” in *Bose-Einstein Condensation in Atomic Gases*, edited by M. Inguscio, S. Stringari, and C. E. Wieman, Proceedings of the International School of Physics “Enrico Fermi”, Course CXL (IOS, Amsterdam, 1999).

References

- [87] A. L. Fetter and A. A. Svidzinsky, “Vortices in a trapped dilute Bose-Einstein condensate,” *J. Phys.: Condens. Matter* **13**, R135 (2001).
- [88] A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-Particle Systems* (McGraw-Hill, New York, 1971).
- [89] R. P. Feynman, “Space-Time Approach to Non-Relativistic Quantum Mechanics,” *Rev. Mod. Phys.* **20**, 367 (1948).
- [90] R. P. Feynman, “Application of quantum mechanics to liquid helium,” in *Progress in Low Temperature Physics*, edited by C. J. Gorter and D. F. Brewer, volume 1 (North-Holland, Amsterdam, 1955).
- [91] R. P. Feynman, “Negative probability,” in *Quantum Implications: Essays in Honour of David Bohm*, edited by B. J. Hiley and F. D. Peat (Routledge, London, 1991).
- [92] R. P. Feynman, *The meaning of it all: thoughts of a citizen scientist* (Addison-Wesley, New York, 1998).
- [93] D. S. Fisher, “Flux-lattice melting in thin-film superconductors,” *Phys. Rev. B* **22**, 1190 (1980).
- [94] M. Flaig and U. R. Fischer, “Global singularity-free solution of the Iordanskii force problem,” *Phys. Rev. B* **74**, 224503 (2006).
- [95] C. Fox, *private communication*.
- [96] M. Frigo and S. G. Johnson, “The Fastest Fourier Transform in the West,” <http://www.fftw.org>.
- [97] M. Frigo and S. G. Johnson, “The Design and Implementation of FFTW3,” *Proc. IEEE* **93**, 216 (2005).
- [98] M. Gajda, “Criterion for Bose-Einstein condensation in a harmonic trap in the case with attractive interactions,” *Phys. Rev. A* **73**, 023603 (2006).
- [99] J. J. García-Ripoll and V. M. Pérez-García, “Anomalous rotational properties of Bose-Einstein condensates in asymmetric traps,” *Phys. Rev. A* **64**, 013602 (2001).
- [100] C. W. Gardiner, “Particle-number-conserving Bogoliubov method which demonstrates the validity of the time-dependent Gross-Pitaevskii equation for a highly condensed Bose gas,” *Phys. Rev. A* **56**, 1414 (1997).

-
- [101] C. W. Gardiner, *Handbook of Stochastic Methods*, 3rd edition (Springer-Verlag, Berlin, 2004).
- [102] C. W. Gardiner, J. R. Anglin, and T. I. A. Fudge, “The stochastic Gross-Pitaevskii equation,” *J. Phys. B* **35**, 1555 (2002).
- [103] C. W. Gardiner and A. S. Bradley, “Analysis of g_2 for the cold collision frequency shift in the hydrogen condensate experiments,” *J. Phys. B* **34**, 4663 (2001).
- [104] C. W. Gardiner and M. J. Davis, “The stochastic Gross-Pitaevskii equation: II,” *J. Phys. B* **36**, 4731 (2003).
- [105] C. W. Gardiner and P. Zoller, “Quantum kinetic theory. III. Quantum kinetic master equation for strongly condensed trapped systems,” *Phys. Rev. A* **58**, 536 (1998).
- [106] C. W. Gardiner and P. Zoller, *Quantum Noise*, 2nd edition (Springer-Verlag, Berlin, 2000).
- [107] S. A. Gardiner and S. A. Morgan, “Number-conserving approach to a minimal self-consistent treatment of condensate and noncondensate dynamics in a degenerate Bose gas,” *Phys. Rev. A* **75**, 043621 (2007).
- [108] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer, and A. Aspect, “Critical Temperature of a Trapped, Weakly Interacting Bose Gas,” *Phys. Rev. Lett.* **92**, 030405 (2004).
- [109] C. Gies, B. P. van Zyl, S. A. Morgan, and D. A. W. Hutchinson, “Finite-temperature theory of the trapped two-dimensional Bose gas,” *Phys. Rev. A* **69**, 023616 (2004).
- [110] S. A. Gifford and G. Baym, “Dislocation-mediated melting in superfluid vortex lattices,” *Phys. Rev. A* **78**, 043607 (2008).
- [111] V. L. Ginzburg and L. P. Pitaevskii, “On the theory of superfluidity,” *Sov. Phys. JETP* **7**, 858 (1958).
- [112] S. Giorgini, “Damping in dilute Bose gases: A mean-field approach,” *Phys. Rev. A* **57**, 2949 (1998).
- [113] S. Giorgini, L. P. Pitaevskii, and S. Stringari, “Condensate fraction and critical temperature of a trapped interacting Bose gas,” *Phys. Rev. A* **54**, R4633 (1996).
- [114] S. Giorgini, L. P. Pitaevskii, and S. Stringari, “Thermodynamics of a Trapped Bose-Condensed Gas,” *J. Low Temp. Phys.* **109**, 309 (1997).

References

- [115] M. Girardeau and R. Arnowitt, “Theory of Many-Boson Systems: Pair Theory,” *Phys. Rev.* **113**, 755 (1959).
- [116] R. J. Glauber, “Optical Coherence and Photon Statistics,” in *Quantum Optics and Electronics*, edited by C. DeWitt, A. Blandin, and C. Cohen-Tannoudji, Proceedings of the Les Houches Summer School of Theoretical Physics, Session XIV (Gordon and Breach, New York, 1965).
- [117] H. Goldstein, *Classical Mechanics* (Addison-Wesley, Cambridge, Massachusetts, 1950).
- [118] R. Graham, “Off-diagonal Long Range Order in Time,” *Nonlinear Phenom. Complex Syst. (Dordrecht, Neth.)* **5**, 349 (2002).
- [119] A. Griffin, “Conserving and gapless approximations for an inhomogeneous Bose gas at finite temperatures,” *Phys. Rev. B* **53**, 9341 (1996).
- [120] D. Guéry-Odelin, “Spinning up and down a Boltzmann gas,” *Phys. Rev. A* **62**, 033607 (2000).
- [121] I. Guillamón, H. Suderow, A. Fernández-Pacheco, J. Sesé, R. Córdoba, J. M. De Teresa, M. R. Ibarra, and S. Vieira, “Direct observation of melting in a two-dimensional superconducting vortex lattice,” *Nature Phys.* **5**, 651 (2009).
- [122] M. Guilleumas and R. Graham, “Off-axis vortices in trapped Bose-condensed gases: Angular momentum and frequency splitting,” *Phys. Rev. A* **64**, 033607 (2001).
- [123] Z. Hadzibabic and J. Dalibard, “Two-dimensional Bose fluids: An atomic physics perspective,” in *Nano optics and atomics: Transport of light and matter waves*, edited by R. Kaiser, D. S. Wiersma, and L. Fallani, Proceedings of the International School of Physics “Enrico Fermi”, Course CLXXIII (2010).
- [124] P. C. Haljan, I. Coddington, P. Engels, and E. A. Cornell, “Driving Bose-Einstein-Condensate Vorticity with a Rotating Normal Cloud,” *Phys. Rev. Lett.* **87**, 210403 (2001).
- [125] H. E. Hall and W. F. Vinen, “The Rotation of Liquid Helium II. II. The Theory of Mutual Friction in Uniformly Rotating Helium II,” *Proc. R. Soc. Lond. A* **238**, 215 (1956).
- [126] M. H. Hayes, *Statistical Digital Signal Processing and Modeling* (Wiley, New York, 1996).

-
- [127] E. A. L. Henn, J. A. Seman, G. Roati, K. M. F. Magalhães, and V. S. Bagnato, “Emergence of Turbulence in an Oscillating Bose-Einstein Condensate,” *Phys. Rev. Lett.* **103**, 045301 (2009).
- [128] G. B. Hess and W. M. Fairbank, “Measurements of Angular Momentum in Superfluid Helium,” *Phys. Rev. Lett.* **19**, 216 (1967).
- [129] M. Hillery, R. F. O’Connell, M. O. Scully, and E. P. Wigner, “Distribution functions in physics: Fundamentals,” *Phys. Rep.* **106**, 121 (1984).
- [130] E. Hodby, G. Hechenblaikner, S. A. Hopkins, O. M. Maragò, and C. J. Foot, “Vortex Nucleation in Bose-Einstein Condensates in an Oblate, Purely Magnetic Potential,” *Phys. Rev. Lett.* **88**, 010405 (2001).
- [131] P. C. Hohenberg and P. C. Martin, “Microscopic theory of superfluid helium,” *Ann. Phys. (N. Y.)* **34**, 291 (1965).
- [132] K. Huang, *Statistical mechanics*, 2nd edition (Wiley, New York, 1987).
- [133] D. A. W. Hutchinson, K. Burnett, R. J. Dodd, S. A. Morgan, M. Rusch, E. Zaremba, N. P. Proukakis, M. Edwards, and C. W. Clark, “Gapless mean-field theory of Bose-Einstein condensates,” *J. Phys. B* **33**, 3825 (2000).
- [134] D. A. W. Hutchinson, R. J. Dodd, and K. Burnett, “Gapless Finite- T Theory of Collective Modes of a Trapped Gas,” *Phys. Rev. Lett.* **81**, 2198 (1998).
- [135] D. A. W. Hutchinson, E. Zaremba, and A. Griffin, “Finite Temperature Excitations of a Trapped Bose Gas,” *Phys. Rev. Lett.* **78**, 1842 (1997).
- [136] S. Inouye, S. Gupta, T. Rosenband, A. P. Chikkatur, A. Görlitz, T. L. Gustavson, A. E. Leanhardt, D. E. Pritchard, and W. Ketterle, “Observation of Vortex Phase Singularities in Bose-Einstein Condensates,” *Phys. Rev. Lett.* **87**, 080402 (2001).
- [137] S. Inouye, T. Pfau, S. Gupta, A. P. Chikkatur, A. Görlitz, D. E. Pritchard, and W. Ketterle, “Phase-coherent amplification of atomic matter waves,” *Nature* **402**, 641 (1999).
- [138] T. Isoshima, J. Huhtamäki, and M. M. Salomaa, “Precessional motion of a vortex in a finite-temperature Bose-Einstein condensate,” *Phys. Rev. A* **69**, 063601 (2004).
- [139] T. Isoshima and K. Machida, “Vortex stabilization in Bose-Einstein condensate of alkali-metal atom gas,” *Phys. Rev. A* **59**, 2203 (1999).

References

- [140] B. Jackson, N. P. Proukakis, C. F. Barenghi, and E. Zaremba, “Finite-temperature vortex dynamics in Bose-Einstein condensates,” *Phys. Rev. A* **79**, 053615 (2009).
- [141] J. Javanainen and J. Ruostekoski, “Symbolic calculation in development of algorithms: split-step methods for the Gross-Pitaevskii equation,” *J. Phys. A* **39**, L179 (2006).
- [142] D. M. Jezek and H. M. Cataldo, “Vortex velocity field in inhomogeneous media: A numerical study in Bose-Einstein condensates,” *Phys. Rev. A* **77**, 043602 (2008).
- [143] D. S. Jin, M. R. Matthews, J. R. Ensher, C. E. Wieman, and E. A. Cornell, “Temperature-Dependent Damping and Frequency Shifts in Collective Excitations of a Dilute Bose-Einstein Condensate,” *Phys. Rev. Lett.* **78**, 764 (1997).
- [144] Yu. Kagan, E. L. Surkov, and G. V. Shlyapnikov, “Evolution of a Bose-condensed gas under variations of the confining potential,” *Phys. Rev. A* **54**, R1753 (1996).
- [145] K. Kasamatsu, M. Tsubota, and M. Ueda, “Nonlinear dynamics of vortex lattice formation in a rotating Bose-Einstein condensate,” *Phys. Rev. A* **67**, 033610 (2003).
- [146] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, “Making, probing and understanding Bose-Einstein condensates,” in *Bose-Einstein Condensation in Atomic Gases*, edited by M. Inguscio, S. Stringari, and C. E. Wieman, Proceedings of the International School of Physics “Enrico Fermi”, Course CXL (IOS, Amsterdam, 1999).
- [147] L. Khaykovich, F. Schreck, G. Ferrari, T. Bourdel, J. Cubizolles, L. D. Carr, Y. Castin, and C. Salomon, “Formation of a Matter-Wave Bright Soliton,” *Science* **296**, 1290 (2002).
- [148] T. W. B. Kibble, “Topology of cosmic domains and strings,” *J. Phys. A* **9**, 1387 (1976).
- [149] W. Kohn, “Cyclotron Resonance and de Haas-van Alphen Oscillations of an Interacting Electron Gas,” *Phys. Rev.* **123**, 1242 (1961).
- [150] S. Komineas, N. R. Cooper, and N. Papanicolaou, “Single vortex states in a confined Bose-Einstein condensate,” *Phys. Rev. A* **72**, 053624 (2005).
- [151] E. Kozik and B. Svistunov, “Kelvin-Wave Cascade and Decay of Superfluid Turbulence,” *Phys. Rev. Lett.* **92**, 035301 (2004).

-
- [152] S. Kragset, E. Babaev, and A. Sudbø, “Effects of boundaries and density inhomogeneity on states of vortex matter in Bose-Einstein condensates at finite temperature,” *Phys. Rev. A* **77**, 043605 (2008).
- [153] M. Krämer, L. Pitaevskii, S. Stringari, and F. Zambelli, “Vortex Nucleation and Quadrupole Deformation of a Rotating Bose-Einstein Condensate,” *Laser Phys.* **12**, 113 (2002).
- [154] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein, and T. Pfau, “The physics of dipolar bosonic quantum gases,” *Rep. Prog. Phys.* **72**, 126401 (2009).
- [155] G. Lamporesi, S. Tung, V. Schweikhard, and E. A. Cornell, “Towards thermal melting of a vortex lattice in a rotating 2D BEC,” poster at ICAP2008, Storrs, Connecticut, U.S.A. (2008).
- [156] L. D. Landau and E. M. Lifshitz, *Mechanics* (Pergamon, Oxford, 1960).
- [157] L. D. Landau and E. M. Lifshitz, *Statistical Physics*, 2nd edition (Pergamon, Oxford, 1969).
- [158] L. D. Landau and E. M. Lifshitz, *Fluid mechanics*, 2nd edition (Pergamon, Oxford, 1987).
- [159] A. E. Leanhardt, A. Görlitz, A. P. Chikkatur, D. Kielpinski, Y. Shin, D. E. Pritchard, and W. Ketterle, “Imprinting Vortices in a Bose-Einstein Condensate using Topological Phases,” *Phys. Rev. Lett.* **89**, 190403 (2002).
- [160] J. L. Lebowitz and O. Penrose, “Modern ergodic theory,” *Phys. Today* **26**(2), 23 (1973).
- [161] A. J. Leggett, “Broken Gauge Symmetry in a Bose Condensate,” in *Bose-Einstein Condensation*, edited by A. Griffin, D. W. Snoke, and S. Stringari (Cambridge University Press, Cambridge, 1995).
- [162] A. J. Leggett, “Superfluidity,” *Rev. Mod. Phys.* **71**, S318 (1999).
- [163] A. J. Leggett, “Topics in the Theory of the Ultracold Dilute Alkali Gases,” in *Bose-Einstein Condensation: From Atomic Physics to Quantum Fluids*, edited by C. M. Savage and M. P. Das, Proceedings of the Thirteenth Physics Summer School (World Scientific, Singapore, 2000).
- [164] A. J. Leggett, “Bose-Einstein condensation in the alkali gases: Some fundamental concepts,” *Rev. Mod. Phys.* **73**, 307 (2001).

References

- [165] M. Lewenstein and L. You, “Quantum Phase Diffusion of a Bose-Einstein Condensate,” *Phys. Rev. Lett.* **77**, 3489 (1996).
- [166] E. M. Lifshitz and L. P. Pitaevskii, *Statistical Physics Part 2: Theory of the Condensed State*, 2nd edition (Pergamon, Oxford, 1980).
- [167] Y.-J. Lin, R. L. Compton, K. Jiménez-García, J. V. Porto, and I. B. Spielman, “Synthetic magnetic fields for ultracold neutral atoms,” *Nature* **462**, 628 (2009).
- [168] M. Linn and A. L. Fetter, “Small-amplitude normal modes of a vortex in a trapped Bose-Einstein condensate,” *Phys. Rev. A* **61**, 063603 (2000).
- [169] C. Lobo and Y. Castin, “Nonclassical scissors mode of a vortex lattice in a Bose-Einstein condensate,” *Phys. Rev. A* **72**, 043606 (2005).
- [170] C. Lobo, A. Sinatra, and Y. Castin, “Vortex Lattice Formation in Bose-Einstein Condensates,” *Phys. Rev. Lett.* **92**, 020403 (2004).
- [171] E. Lundh, J.-P. Martikainen, and K.-A. Suominen, “Vortex nucleation in Bose-Einstein condensates in time-dependent traps,” *Phys. Rev. A* **67**, 063604 (2003).
- [172] E. Lundh, C. J. Pethick, and H. Smith, “Zero-temperature properties of a trapped Bose-condensed gas: Beyond the Thomas-Fermi approximation,” *Phys. Rev. A* **55**, 2126 (1997).
- [173] E. Lundh, C. J. Pethick, and H. Smith, “Vortices in Bose-Einstein-condensed atomic clouds,” *Phys. Rev. A* **58**, 4816 (1998).
- [174] K. W. Madison, F. Chevy, V. Bretin, and J. Dalibard, “Stationary States of a Rotating Bose-Einstein Condensate: Routes to Vortex Nucleation,” *Phys. Rev. Lett.* **86**, 4443 (2001).
- [175] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard, “Vortex Formation in a Stirred Bose-Einstein Condensate,” *Phys. Rev. Lett.* **84**, 806 (2000).
- [176] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard, “Vortices in a stirred Bose-Einstein condensate,” *J. Mod. Opt.* **47**, 2715 (2000).
- [177] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, Cambridge, 1995).

-
- [178] O. M. Maragò, S. A. Hopkins, J. Arlt, E. Hodby, G. Hechenblaikner, and C. J. Foot, “Observation of the Scissors Mode and Evidence for Superfluidity of a Trapped Bose-Einstein Condensed Gas,” *Phys. Rev. Lett.* **84**, 2056 (2000).
- [179] MATLAB, The Mathworks Inc., <http://www.mathworks.com>.
- [180] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman, and E. A. Cornell, “Vortices in a Bose-Einstein Condensate,” *Phys. Rev. Lett.* **83**, 2498 (1999).
- [181] T. Mizushima, Y. Kawaguchi, K. Machida, T. Ohmi, T. Isoshima, and M. M. Salomaa, “Collective Oscillations of Vortex Lattices in Rotating Bose-Einstein Condensates,” *Phys. Rev. Lett.* **92**, 060407 (2004).
- [182] C. Mora and Y. Castin, “Extension of Bogoliubov theory to quasicondensates,” *Phys. Rev. A* **67**, 053615 (2003).
- [183] S. A. Morgan, “A gapless theory of Bose-Einstein condensation in dilute gases at finite temperature,” *J. Phys. B* **33**, 3847 (2000).
- [184] S. A. Morgan, S. Choi, K. Burnett, and M. Edwards, “Nonlinear mixing of quasiparticles in an inhomogeneous Bose condensate,” *Phys. Rev. A* **57**, 3818 (1998).
- [185] O. Morsch and M. Oberthaler, “Dynamics of Bose-Einstein condensates in optical lattices,” *Rev. Mod. Phys.* **78**, 179 (2006).
- [186] M. Naraschewski and R. J. Glauber, “Spatial coherence and density correlations of trapped Bose gases,” *Phys. Rev. A* **59**, 4595 (1999).
- [187] M. Naraschewski and D. M. Stamper-Kurn, “Analytical description of a trapped semi-ideal Bose gas at finite temperature,” *Phys. Rev. A* **58**, 2423 (1998).
- [188] H. M. Nilsen, G. Baym, and C. J. Pethick, “Velocity of vortices in inhomogeneous Bose-Einstein condensates,” *Proc. Natl. Acad. Sci. U.S.A.* **103**, 7978 (2006).
- [189] A. A. Norrie, *A Classical Field Treatment of Colliding Bose-Einstein Condensates*, Ph.D. thesis, University of Otago (2005).
- [190] A. A. Norrie, R. J. Ballagh, and C. W. Gardiner, “Quantum turbulence and correlations in Bose-Einstein condensate collisions,” *Phys. Rev. A* **73**, 043617 (2006).
- [191] N. Papanicolaou, S. Komineas, and N. R. Cooper, “Virial theorems for vortex states in a confined Bose-Einstein condensate,” *Phys. Rev. A* **72**, 053609 (2005).

References

- [192] N. G. Parker, *Numerical Studies of Vortices and Dark Solitons in Atomic Bose-Einstein Condensates*, Ph.D. thesis, University of Durham (2004).
- [193] N. G. Parker and C. S. Adams, *private communication*.
- [194] N. G. Parker and C. S. Adams, “Emergence and Decay of Turbulence in Stirred Atomic Bose-Einstein Condensates,” *Phys. Rev. Lett.* **95**, 145301 (2005).
- [195] N. G. Parker and C. S. Adams, “Response of an atomic Bose-Einstein condensate to a rotating elliptical trap,” *J. Phys. B* **39**, 43 (2006).
- [196] N. G. Parker, R. M. W. van Bijnen, and A. M. Martin, “Instabilities leading to vortex lattice formation in rotating Bose-Einstein condensates,” *Phys. Rev. A* **73**, 061603(R) (2006).
- [197] R. K. Pathria, *Statistical Mechanics*, 2nd edition (Butterworth-Heinemann, Oxford, 1996).
- [198] W. Pauli, “The Connection Between Spin and Statistics,” *Phys. Rev.* **58**, 716 (1940).
- [199] A. A. Penckwitt, *Rotating Bose-Einstein Condensates*, Ph.D. thesis, University of Otago (2003).
- [200] A. A. Penckwitt, R. J. Ballagh, and C. W. Gardiner, “Nucleation, Growth, and Stabilization of Bose-Einstein Condensate Vortex Lattices,” *Phys. Rev. Lett.* **89**, 260402 (2002).
- [201] O. Penrose and L. Onsager, “Bose-Einstein Condensation and Liquid Helium,” *Phys. Rev.* **104**, 576 (1956).
- [202] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Westview, Boulder, Colorado, 1995).
- [203] C. J. Pethick and L. P. Pitaevskii, “Criterion for Bose-Einstein condensation for particles in traps,” *Phys. Rev. A* **62**, 033609 (2000).
- [204] C. J. Pethick and H. Smith, *Bose-Einstein Condensation in Dilute Gases* (Cambridge University Press, Cambridge, 2002).
- [205] D. S. Petrov, M. Holzmann, and G. V. Shlyapnikov, “Bose-Einstein Condensation in Quasi-2D Trapped Gases,” *Phys. Rev. Lett.* **84**, 2551 (2000).
- [206] L. P. Pitaevskii and S. Stringari, “Landau damping in dilute Bose gases,” *Phys. Lett. A* **235**, 398 (1997).

-
- [207] L. P. Pitaevskii and S. Stringari, *Bose-Einstein Condensation* (Clarendon Press, Oxford, 2003).
- [208] W. V. Pogosov and K. Machida, “Thermal fluctuations of vortex clusters in quasi-two-dimensional Bose-Einstein condensates,” *Phys. Rev. A* **74**, 023622 (2006).
- [209] A. Polkovnikov, *private communication*.
- [210] A. Polkovnikov, “Quantum corrections to the dynamics of interacting bosons: Beyond the truncated Wigner approximation,” *Phys. Rev. A* **68**, 053604 (2003).
- [211] A. Polkovnikov, “Phase space representation of quantum dynamics,” *Ann. Phys. (N.Y.)* **325**, 1790 (2010).
- [212] A. Polkovnikov and D.-W. Wang, “Effect of Quantum Fluctuations on the Dipolar Motion of Bose-Einstein Condensates in Optical Lattices,” *Phys. Rev. Lett.* **93**, 070401 (2004).
- [213] V. N. Popov, *Functional Integrals in Quantum Field Theory and Statistical Physics* (Kluwer Academic, Dordrecht, 1983).
- [214] A. Posazhennikova, “Colloquium: Weakly interacting, dilute Bose gases in 2D,” *Rev. Mod. Phys.* **78**, 1111 (2006).
- [215] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd edition (Cambridge University Press, Cambridge, 1992).
- [216] N. P. Proukakis and K. Burnett, “Generalized Mean Fields for Trapped Atomic Bose-Einstein Condensates,” *J. Res. Natl. Inst. Stand. Technol.* **101**, 457 (1996).
- [217] N. P. Proukakis and B. Jackson, “Finite-temperature models of Bose-Einstein condensation,” *J. Phys. B* **41**, 203002 (2008).
- [218] N. P. Proukakis, S. A. Morgan, S. Choi, and K. Burnett, “Comparison of gapless mean-field theories for trapped Bose-Einstein condensates,” *Phys. Rev. A* **58**, 2435 (1998).
- [219] H. Pu, C. K. Law, J. H. Eberly, and N. P. Bigelow, “Coherent disintegration and stability of vortices in trapped Bose condensates,” *Phys. Rev. A* **59**, 1533 (1999).
- [220] A. Recati, F. Zambelli, and S. Stringari, “Overcritical Rotation of a Trapped Bose-Einstein Condensate,” *Phys. Rev. Lett.* **86**, 377 (2001).

References

- [221] F. Reif, *Fundamentals of Statistical and Thermal Physics* (McGraw-Hill, New York, 1965).
- [222] J. W. Robinson, *Atomic Spectroscopy*, 2nd edition (Marcel Dekker, New York, 1996).
- [223] D. S. Rokhsar, “Vortex Stability and Persistent Currents in Trapped Bose Gases,” *Phys. Rev. Lett.* **79**, 2164 (1997).
- [224] H. H. Rugh, “Dynamical Approach to Temperature,” *Phys. Rev. Lett.* **78**, 772 (1997).
- [225] R. Z. Sagdeev, D. A. Usikov, and G. M. Zaslavsky, “Nonlinear Physics: From the Pendulum to Turbulence and Chaos,” (Harwood Academic, Chur, 1988).
- [226] D. R. Scherer, C. N. Weiler, T. W. Neely, and B. P. Anderson, “Vortex Formation by Merging of Multiple Trapped Bose-Einstein Condensates,” *Phys. Rev. Lett.* **98**, 110402 (2007).
- [227] H. Schmidt, K. Góral, F. Floegel, M. Gajda, and K. Rzążewski, “Probing the classical field approximation – thermodynamics and decaying vortices,” *J. Opt. B* **5**, S96 (2003).
- [228] B. I. Schneider and D. L. Feder, “Numerical approach to the ground and excited states of a Bose-Einstein condensed gas confined in a completely anisotropic trap,” *Phys. Rev. A* **59**, 2232 (1999).
- [229] S. S. Schweber, *An Introduction to Relativistic Quantum Field Theory* (Harper and Row, New York, 1961).
- [230] M. O. Scully and M. S. Zubairy, *Quantum Optics* (Cambridge University Press, Cambridge, 1997).
- [231] R. Seiringer, “Vortices and spontaneous symmetry breaking in rotating Bose gases,” in *Mathematical Results in Quantum Mechanics: Proceedings of the QMath10 Conference*, edited by I. Beltita, G. Nenciu, and R. Purice (World Scientific, Singapore, 2008).
- [232] J. P. Sethna, *Statistical Mechanics: Entropy, Order Parameters and Complexity* (Oxford University Press, Oxford, 2006).
- [233] D. E. Sheehy and L. Radzihovsky, “Vortex lattice inhomogeneity in spatially inhomogeneous superfluids,” *Phys. Rev. A* **70**, 051602(R) (2004).

-
- [234] D. E. Sheehy and L. Radzihovsky, “Vortices in spatially inhomogeneous superfluids,” *Phys. Rev. A* **70**, 063620 (2004).
- [235] A. Sinatra and Y. Castin, “Genuine phase diffusion of a Bose-Einstein condensate in the microcanonical ensemble: A classical field study,” *Phys. Rev. A* **78**, 053615 (2008).
- [236] A. Sinatra, Y. Castin, and C. Lobo, “A Monte Carlo formulation of the Bogolubov theory,” *J. Mod. Opt.* **47**, 2629 (2000).
- [237] A. Sinatra, P. O. Fedichev, Y. Castin, J. Dalibard, and G. V. Shlyapnikov, “Dynamics of Two Interacting Bose-Einstein Condensates,” *Phys. Rev. Lett.* **82**, 251 (1999).
- [238] A. Sinatra, C. Lobo, and Y. Castin, “Classical-Field Method for Time Dependent Bose-Einstein Condensed Gases,” *Phys. Rev. Lett.* **87**, 210404 (2001).
- [239] A. Sinatra, C. Lobo, and Y. Castin, “The truncated Wigner method for Bose-condensed gases: limits of validity and applications,” *J. Phys. B* **35**, 3599 (2002).
- [240] S. Sinha and Y. Castin, “Dynamic Instability of a Rotating Bose-Einstein Condensate,” *Phys. Rev. Lett.* **87**, 190402 (2001).
- [241] D. V. Skryabin, “Instabilities of vortices in a binary mixture of trapped Bose-Einstein condensates: Role of collective excitations with positive and negative energies,” *Phys. Rev. A* **63**, 013602 (2000).
- [242] M. Snoek and H. T. C. Stoof, “Theory of vortex-lattice melting in a one-dimensional optical lattice,” *Phys. Rev. A* **74**, 033615 (2006).
- [243] E. B. Sonin, “Vortex oscillations and hydrodynamics of rotating superfluids,” *Rev. Mod. Phys.* **59**, 87 (1987).
- [244] E. B. Sonin, “Magnus force in superfluids and superconductors,” *Phys. Rev. B* **55**, 485 (1997).
- [245] R. A. Sorensen, “Nuclear Moment of Inertia at High Spin,” *Rev. Mod. Phys.* **45**, 353 (1973).
- [246] M. J. Steel, M. K. Olsen, L. I. Plimak, P. D. Drummond, S. M. Tan, M. J. Collett, D. F. Walls, and R. Graham, “Dynamical quantum noise in trapped Bose-Einstein condensates,” *Phys. Rev. A* **58**, 4824 (1998).

References

- [247] S. Stock, Z. Hadzibabic, B. Battelier, M. Cheneau, and J. Dalibard, “Observation of Phase Defects in Quasi-Two-Dimensional Bose-Einstein Condensates,” *Phys. Rev. Lett.* **95**, 190403 (2005).
- [248] H. T. C. Stoof, “Field Theory for Trapped Atomic Gases,” in *Coherent Atomic Matter Waves*, edited by R. Kaiser, C. Westbrook, and F. David, Proceedings of the Les Houches Summer School of Theoretical Physics, Session LXXII (EDP Sciences and Springer-Verlag, Berlin, 2001).
- [249] S. Stringari, “Collective Excitations of a Trapped Bose-Condensed Gas,” *Phys. Rev. Lett.* **77**, 2360 (1996).
- [250] S. Stringari, “Moment of Inertia and Superfluidity of a Trapped Bose Gas,” *Phys. Rev. Lett.* **76**, 1405 (1996).
- [251] S. Stringari, “Phase Diagram of Quantized Vortices in a Trapped Bose-Einstein Condensed Gas,” *Phys. Rev. Lett.* **82**, 4371 (1999).
- [252] E. C. G. Sudarshan and N. Mukunda, *Classical Dynamics: A Modern Perspective* (Wiley, New York, 1974).
- [253] V. K. Tkachenko, “On vortex lattices,” *Sov. Phys. JETP* **22**, 1282 (1966).
- [254] M. Tsubota, K. Kasamatsu, and M. Ueda, “Vortex lattice formation in a rotating Bose-Einstein condensate,” *Phys. Rev. A* **65**, 023603 (2002).
- [255] S. Tung, V. Schweikhard, and E. A. Cornell, “Observation of Vortex Pinning in Bose-Einstein Condensates,” *Phys. Rev. Lett.* **97**, 240402 (2006).
- [256] S. Viefers, “Quantum Hall physics in rotating Bose-Einstein condensates,” *J. Phys.: Condens. Matter* **20**, 123202 (2008).
- [257] S. M. M. Virtanen, T. P. Simula, and M. M. Salomaa, “Comparison of mean-field theories for vortices in trapped Bose-Einstein condensates,” *J. Phys.: Condens. Matter* **13**, L819 (2001).
- [258] O. K. Vorov, P. Van Isacker, M. S. Hussein, and K. Bartschat, “Nucleation and Growth of Vortices in a Rotating Bose-Einstein Condensate,” *Phys. Rev. Lett.* **95**, 230406 (2005).
- [259] S. Wang, Y. A. Sergeev, C. F. Barenghi, and M. A. Harrison, “Two-Particle Separation in the Point Vortex Gas Model of Superfluid Turbulence,” *J. Low Temp. Phys.* **149**, 65 (2007).

-
- [260] C. N. Weiler, T. W. Neely, D. R. Scherer, A. S. Bradley, M. J. Davis, and B. P. Anderson, “Spontaneous vortices in the formation of Bose-Einstein condensates,” *Nature* **455**, 948 (2008).
- [261] S. Weinberg, *The Quantum Theory of Fields*, volume 1 (Cambridge University Press, Cambridge, 1995).
- [262] G. C. Wick, “On the Magnetic Field of a Rotating Molecule,” *Phys. Rev.* **73**, 51 (1948).
- [263] E. Wigner, “On the Quantum Correction For Thermodynamic Equilibrium,” *Phys. Rev.* **40**, 749 (1932).
- [264] B. G. Wild and D. A. W. Hutchinson, “Precession of vortices in dilute Bose-Einstein condensates at finite temperature,” *Phys. Rev. A* **80**, 035603 (2009).
- [265] N. K. Wilkin, J. M. F. Gunn, and R. A. Smith, “Do Attractive Bosons Condense?” *Phys. Rev. Lett.* **80**, 2265 (1998).
- [266] J. E. Williams, E. Zaremba, B. Jackson, T. Nikuni, and A. Griffin, “Dynamical Instability of a Condensate Induced by a Rotating Thermal Gas,” *Phys. Rev. Lett.* **88**, 070401 (2002).
- [267] T. M. Wright, R. J. Ballagh, A. S. Bradley, P. B. Blakie, and C. W. Gardiner, “Dynamical thermalization and vortex formation in stirred two-dimensional Bose-Einstein condensates,” *Phys. Rev. A* **78**, 063601 (2008).
- [268] T. M. Wright, P. B. Blakie, and R. J. Ballagh, “Temporal coherence, anomalous moments, and pairing correlations in the classical-field description of a degenerate Bose gas,” *Phys. Rev. A* **82**, 013621 (2010).
- [269] T. M. Wright, A. S. Bradley, and R. J. Ballagh, “Finite-temperature dynamics of a single vortex in a Bose-Einstein condensate: Equilibrium precession and rotational symmetry breaking,” *Phys. Rev. A* **80**, 053624 (2009).
- [270] T. M. Wright, A. S. Bradley, and R. J. Ballagh, “Nonequilibrium dynamics of vortex arrest in a finite-temperature Bose-Einstein condensate,” *Phys. Rev. A* **81**, 013610 (2010).
- [271] T. Yamada, Y. Funaki, H. Horiuchi, G. Röpke, P. Schuck, and A. Tohsaki, “Internal one-particle density matrix for Bose-Einstein condensates with finite number of particles in a harmonic potential,” *Phys. Rev. C* **79**, 054314 (2009).

References

- [272] V. E. Zakharov, V. S. L'vov, and G. Falkovich, *Kolmogorov Spectra of Turbulence 1: Wave Turbulence* (Springer-Verlag, Berlin, 1992).
- [273] F. Zambelli and S. Stringari, “Moment of inertia and quadrupole response function of a trapped superfluid,” *Phys. Rev. A* **63**, 033602 (2001).
- [274] E. Zaremba, T. Nikuni, and A. Griffin, “Dynamics of Trapped Bose Gases at Finite Temperatures,” *J. Low Temp. Phys.* **116**, 277 (1999).
- [275] A. Zee, *Quantum Field Theory in a Nutshell* (Princeton University Press, Princeton, New Jersey, 2003).
- [276] O. N. Zhuravlev, A. E. Muryshev, and P. O. Fedichev, “Dissipative dynamics of vortex arrays in anisotropic traps,” *Phys. Rev. A* **64**, 053601 (2001).
- [277] W. H. Zurek, “Cosmological experiments in superfluid helium?” *Nature* **317**, 505 (1985).
- [278] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck, and W. Ketterle, “Vortices and superfluidity in a strongly interacting Fermi gas,” *Nature* **435**, 1047 (2005).